Prospects for Petascale Climate Modeling:
Can Kilo-Processors, Exa-Flops and Peta-Bytes * Make a Difference in Predicting Earth's Climate?

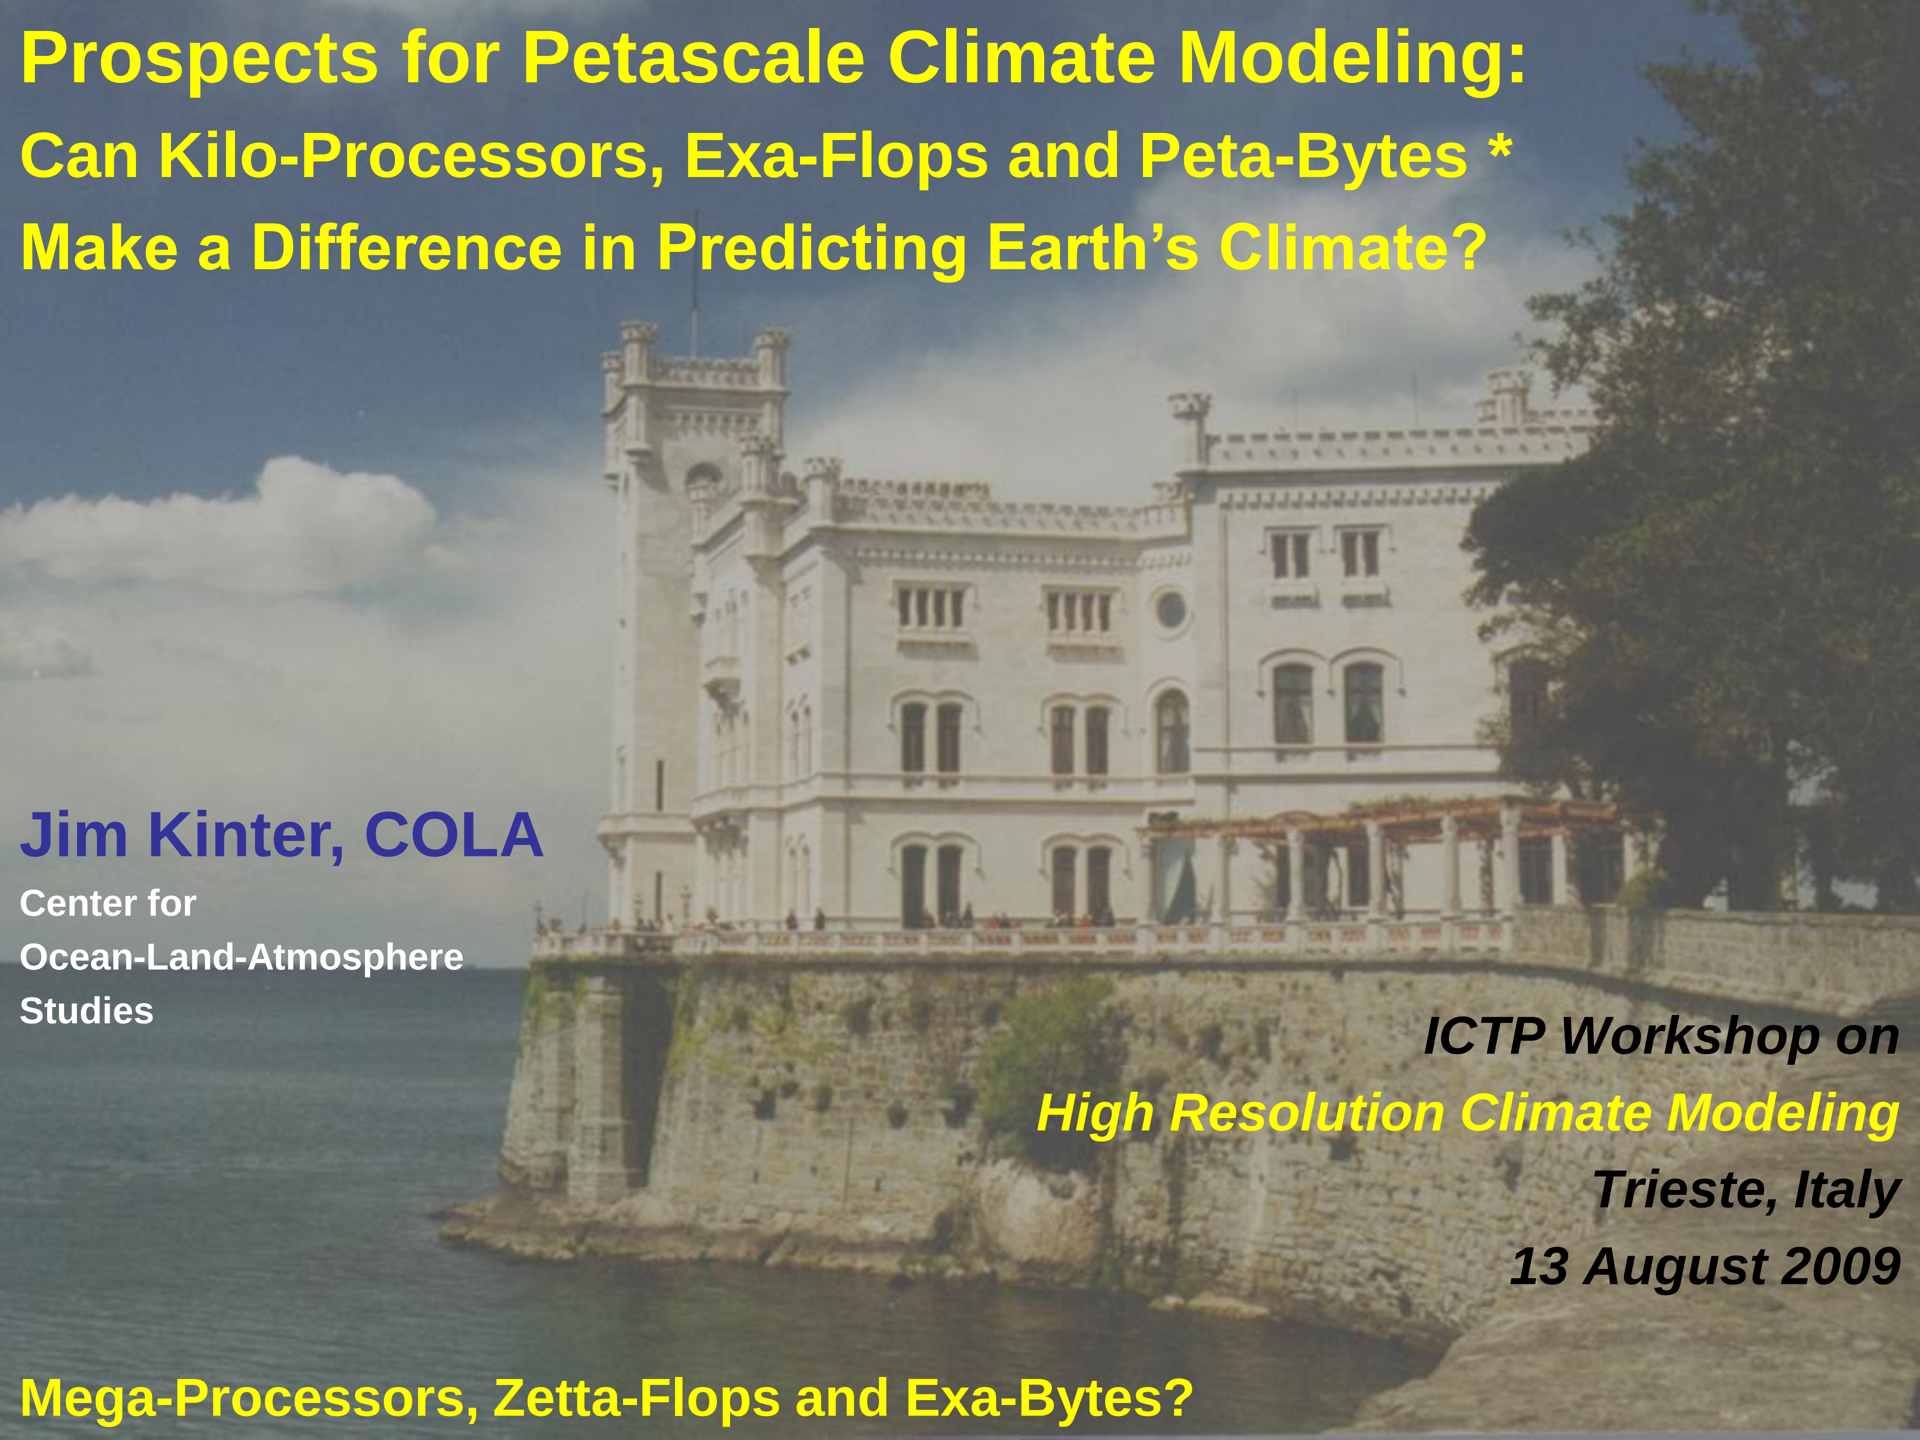

Jim Kinter, COLA

Center for Ocean-Land-Atmosphere Studies

ICTP Workshop on
High Resolution Climate Modeling
Trieste, Italy
13 August 2009

Mega-Processors, Zetta-Flops and Exa-Bytes?

Prefix	Symbol	Power of 10	Power of 2 (-byte)	
yocto-	y	10^{-24}	--	
zepto-	z	10^{-21}	--	
atto-	a	10^{-18}	--	
femto-	f	10^{-15}	--	
pico-	p	10^{-12}	--	
nano-	n	10^{-9}	--	
micro-	m	10^{-6}	--	
milli-	m	10^{-3}	--	
centi-	c	10^{-2}	--	
deci-	d	10^{-1}	--	
	--	10^{0}	--	
deka-	D	10^{1}	--	
hecto-	h	10^{2}	--	
kilo-	**k or K**	$\mathbf{10^{3}}$	$\mathbf{2^{10}}$	**processors (~100K-cores)**
mega-	M	10^{6}	2^{20}	
giga-	G	10^{9}	2^{30}	
tera-	T	10^{12}	2^{40}	
peta-	**P**	$\mathbf{10^{15}}$	$\mathbf{2^{50}}$	**bytes**
exa-	**E**	$\mathbf{10^{18}}$	$\mathbf{2^{60}}$	**flops ***
zetta-	Z	10^{21}	2^{70}	
yotta-	Y	10^{24}	2^{80}	

* e.g. Cray XT5 (kraken) at NICS: 607 peak teraflops * 86400 sec/day = 50 (peak) exaflops/day

10 Points

1. Climate change is most difficult and important problem
2. Scientific consensus: humans contribute to climate change
3. World, urgently engaged in global mitigation and regional adaptation strategies, needs detailed regional information
4. IPCC AR4: considerable uncertainty in predictions of magnitude of global change
5. … and uncertainties in regional climate are even bigger
6. Climate prediction is very computationally demanding
7. Climate spatial scales span 10 decades, but current models resolve less than 4 decades
8. Within 6 years, peak capability of 100 petaflops, 10^7 computing units
9. Parallel computation means new software, algorithms and models
10. … and new methods in workflow management, data management, and visualization

Climate Change

1. Climate change is arguably the most difficult and most important problem facing governments, industries and societies across the Earth.

CO_2

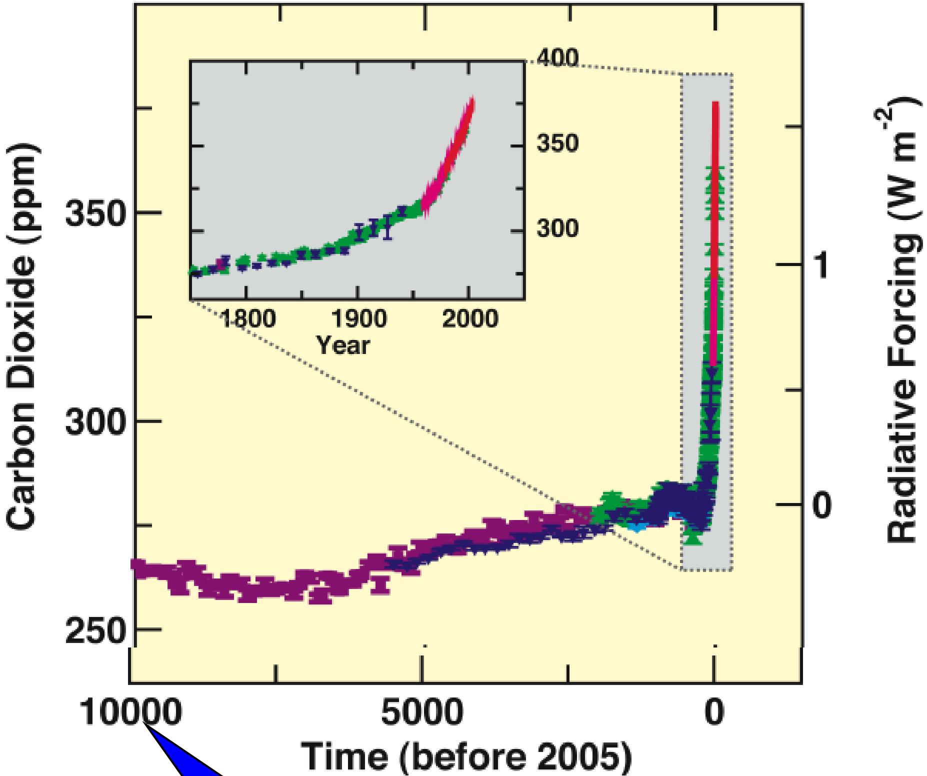

end of last ice age

CH_4

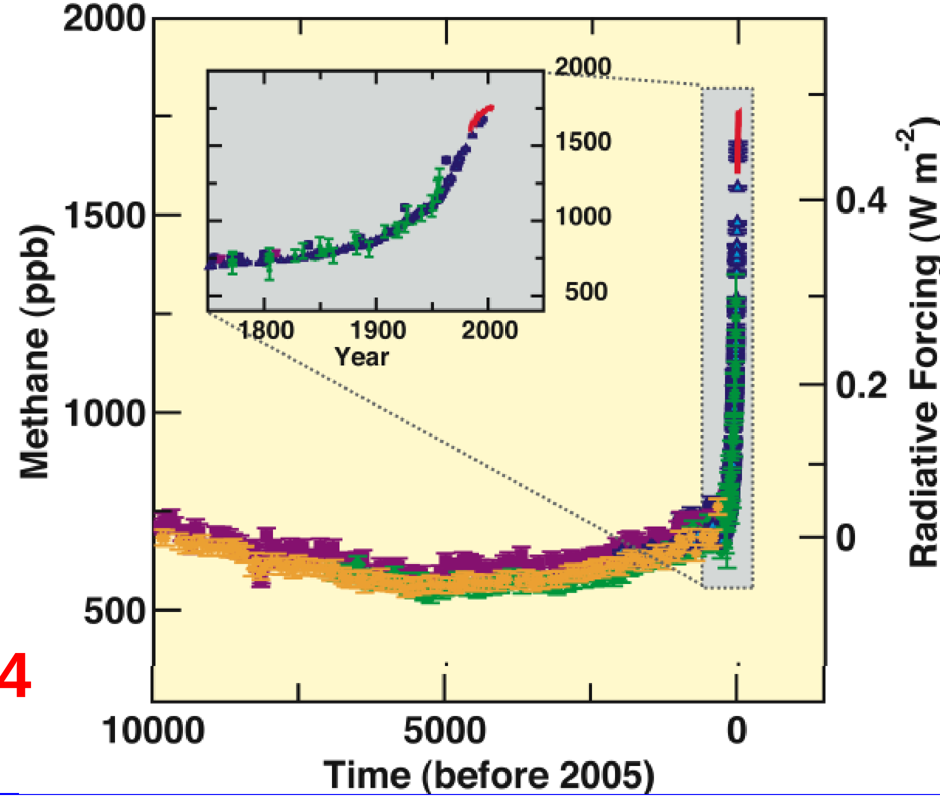

Cost of Climate Change

- Consequences of climate change estimated at several % of world GDP by mid- to late-21st century
- Mitigation efforts to reduce impacts of climate change estimated at 1% of world GDP

(world GDP in 2007: €34 trillion)

Stern Report, 2007

Humans $\Rightarrow$ Climate Change

2. A scientific consensus has emerged that human activities are contributing to climate change.

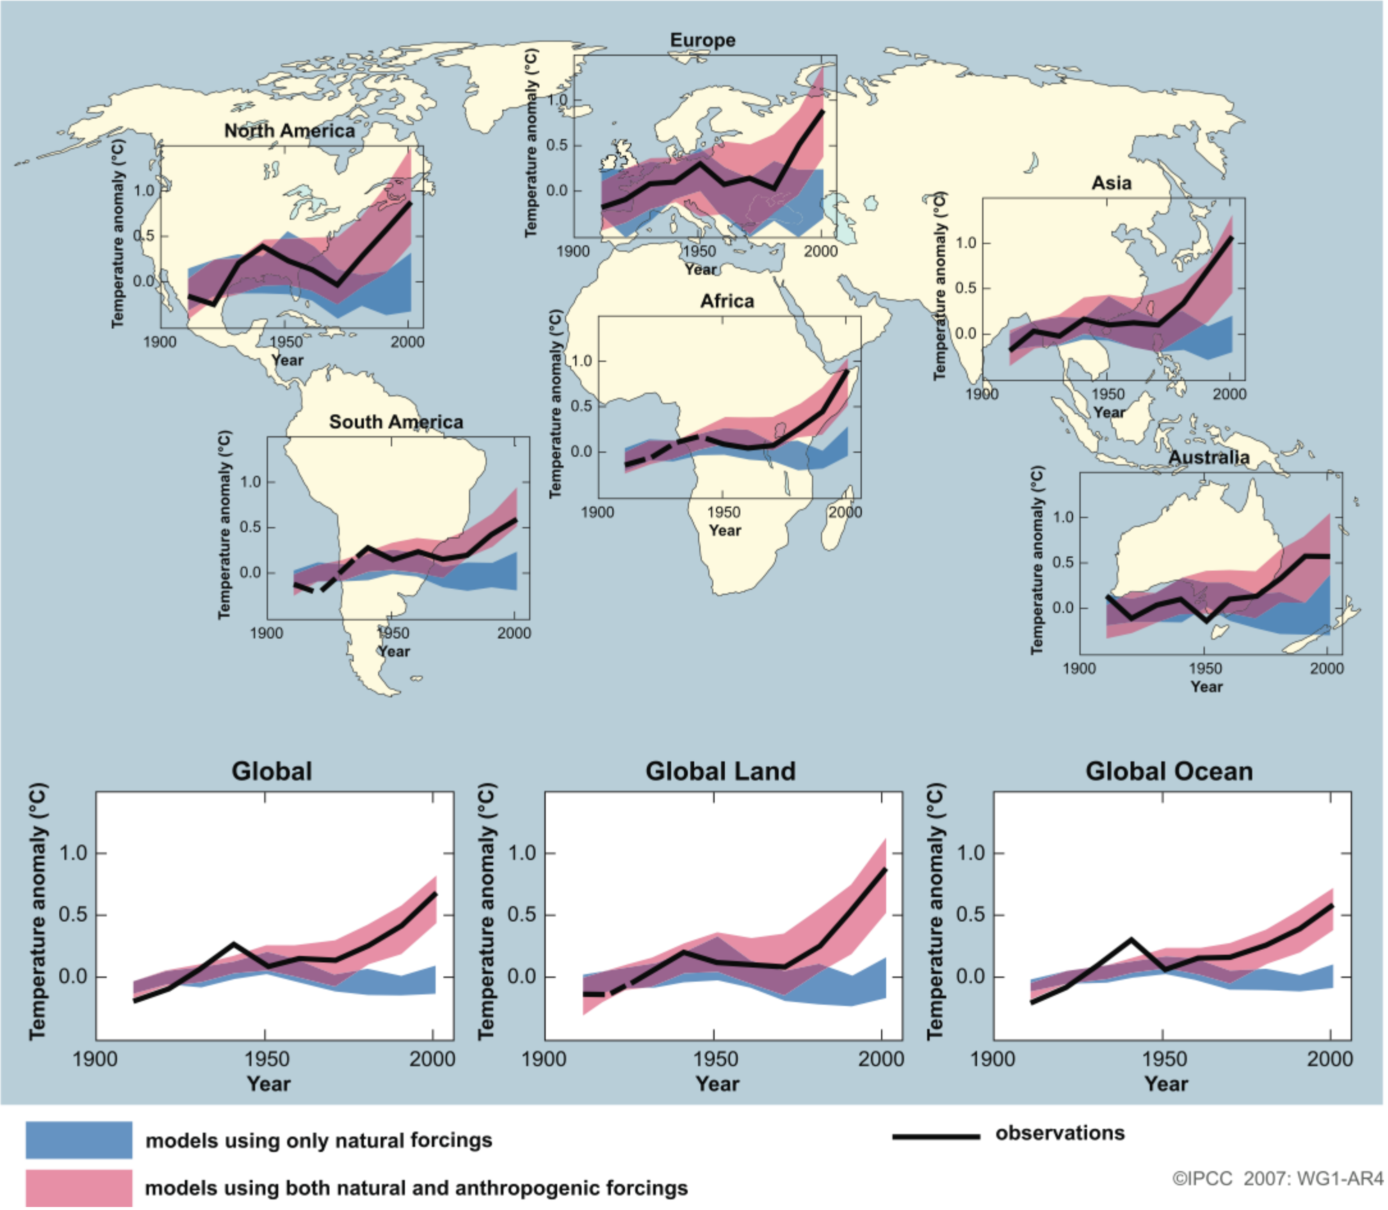
Europe
North America
Asia
Africa
South America
Australia
Global
Global Land
Global Ocean
Temperature anomaly (°C)
1.0
0.5
0.0
1900
1950
2000
Year
models using only natural forcings
models using both natural and anthropogenic forcings
observations

Global Mitigation; Regional Adaptation

3. The world community is now urgently engaged in discussion of strategies for global mitigation of, and regional adaptation to, climate change, which require detailed regional information such as is provided by climate models.

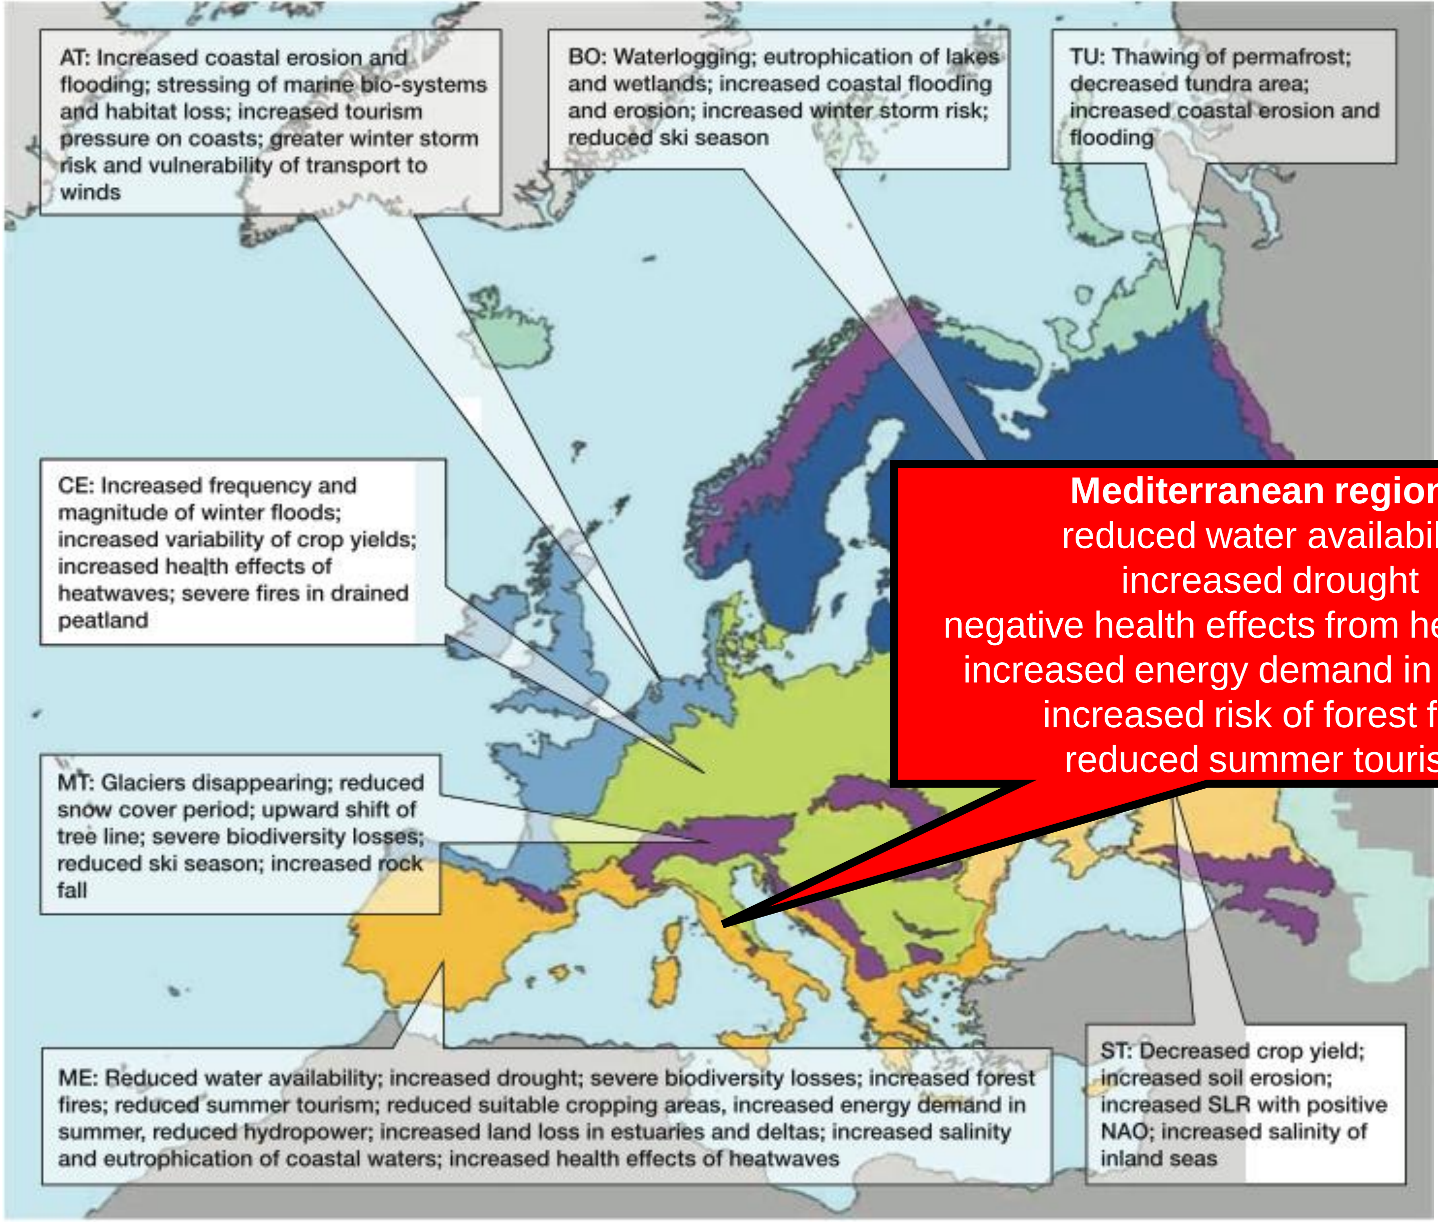

Figure 12.3. *Key vulnerabilities of European systems and sectors to climate change during the 21st century for the main biogeographic regions of Europe (EEA, 2004a): TU: Tundra, pale turquoise. BO: Boreal, dark blue. AT: Atlantic, light blue. CE: Central, green; includes the Pannonian Region. MT: Mountains, purple. ME: Mediterranean, orange; includes the Black Sea region. ST: Steppe, cream. SLR: sea-level rise. NAO: North Atlantic Oscillation. Copyright EEA, Copenhagen. http://www.eea.europa.eu*

Summer 2003 European Heat Wave

How likely is such an event in the future?

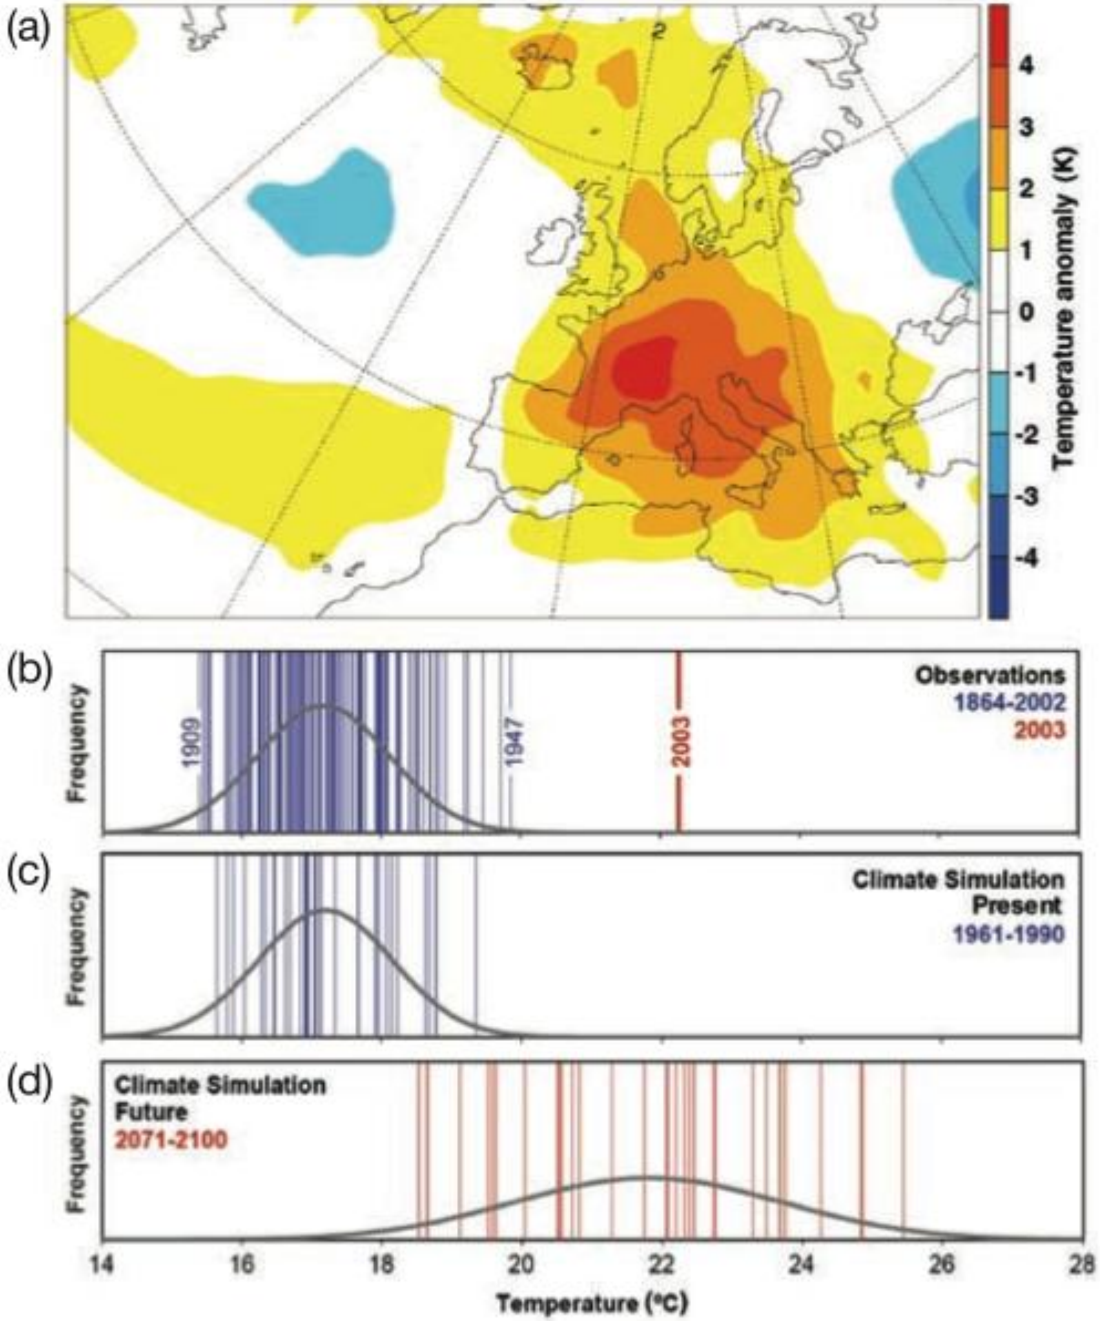

Figure 12.4. *Characteristics of the summer 2003 heatwave (adapted from Schär et al., 2004). (a) JJA temperature anomaly with respect to 1961 to 1990. (b) to (d): JJA temperatures for Switzerland observed during 1864 to 2003 (b), simulated using a regional climate model for the period 1961 to 1990 (c) and simulated for 2071 to 2100 under the A2 scenario using boundary data from the HadAM3H GCM (d). In panels (b) to (d): the black line shows the theoretical frequency distribution of mean summer temperature for the time-period considered, and the vertical blue and red bars show the mean summer temperature for individual years. Reprinted by permission from Macmillan Publishers Ltd. [Nature] (Schär et al., 2004), copyright 2004.*

Uncertainty

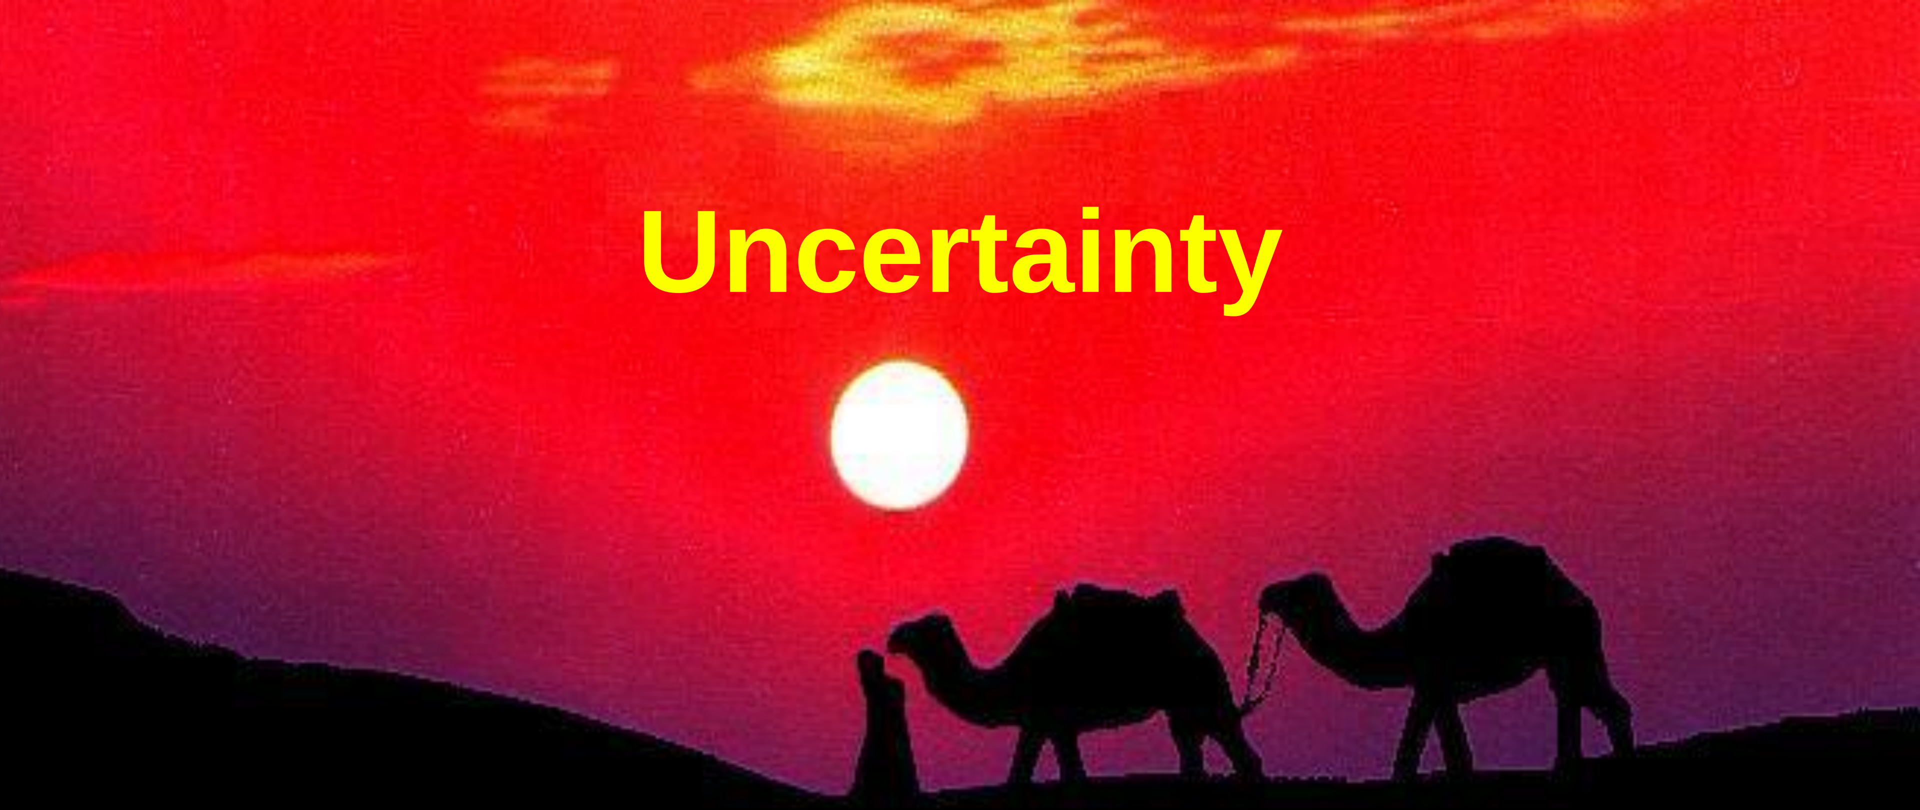

4. As reflected in the IPCC Fourth Assessment Report, there continues to be considerable uncertainty in predictions of the magnitude of global warming.

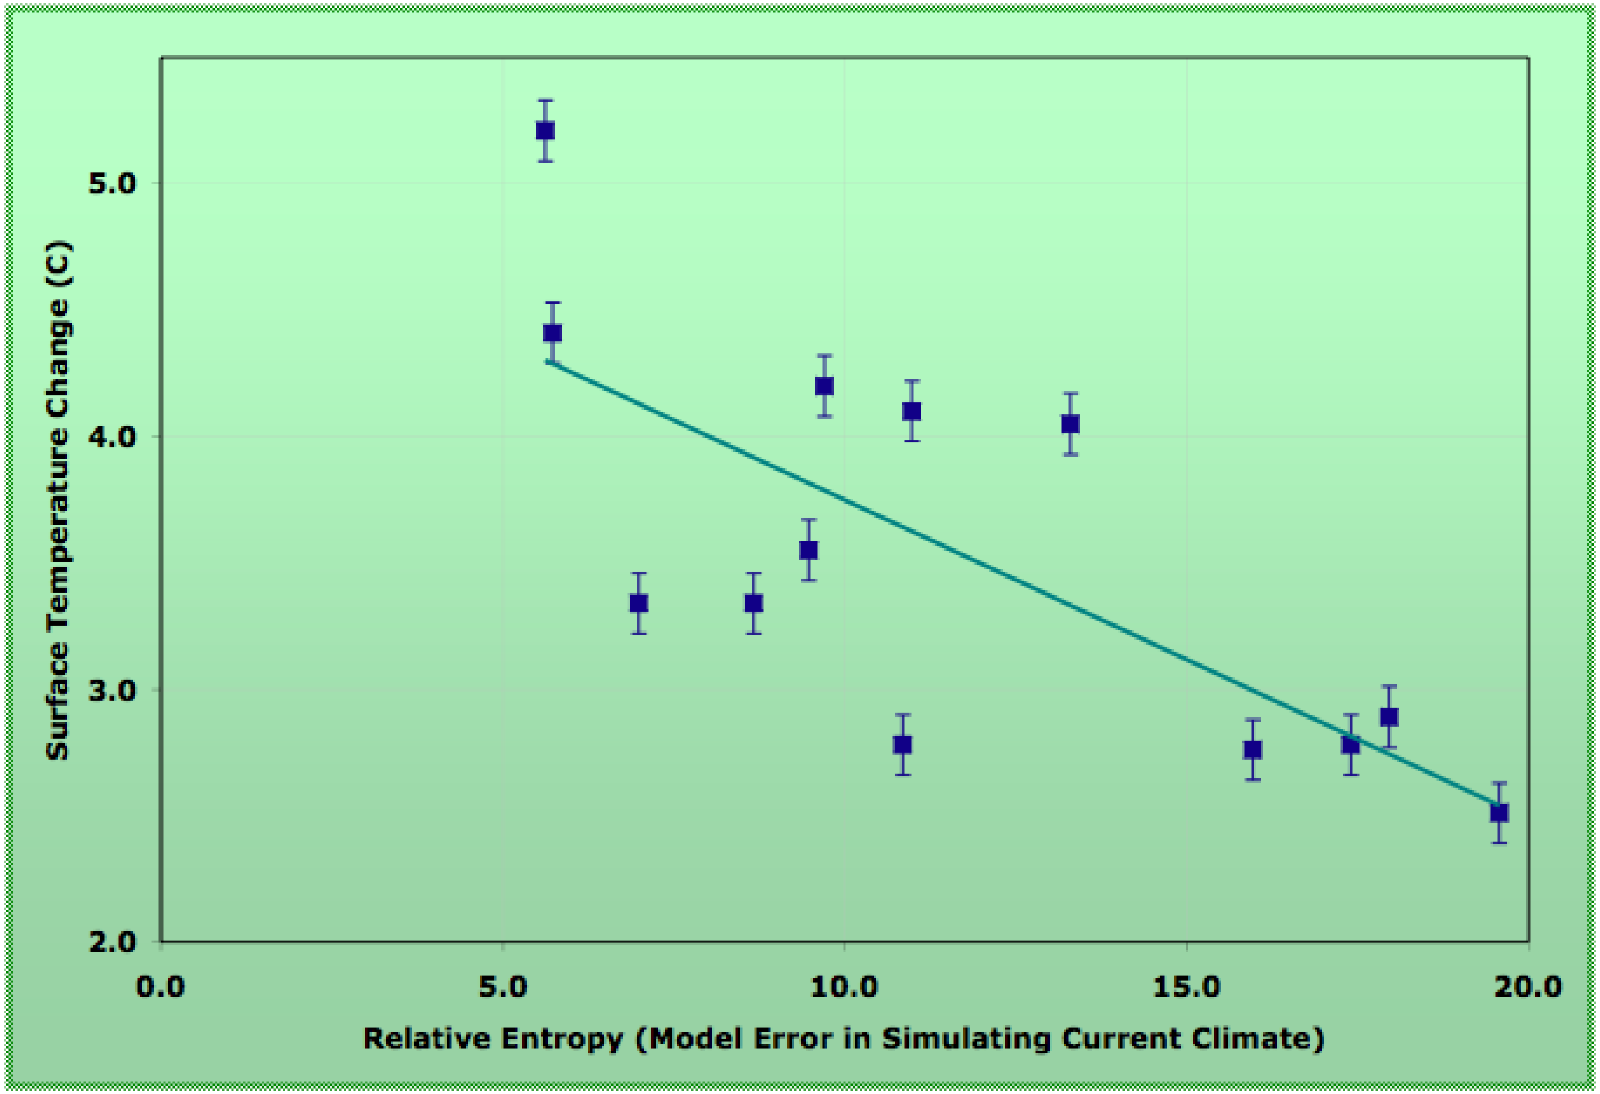

Climate Model Fidelity and Projections of Climate Change

J. Shukla, T. DelSole, M. Fennessy, J. Kinter and D. Paolino

Geophys. Research Letters, **33**, doi10.1029/2005GL025579, 2006

Clouds: Still the Largest Source of Uncertainty

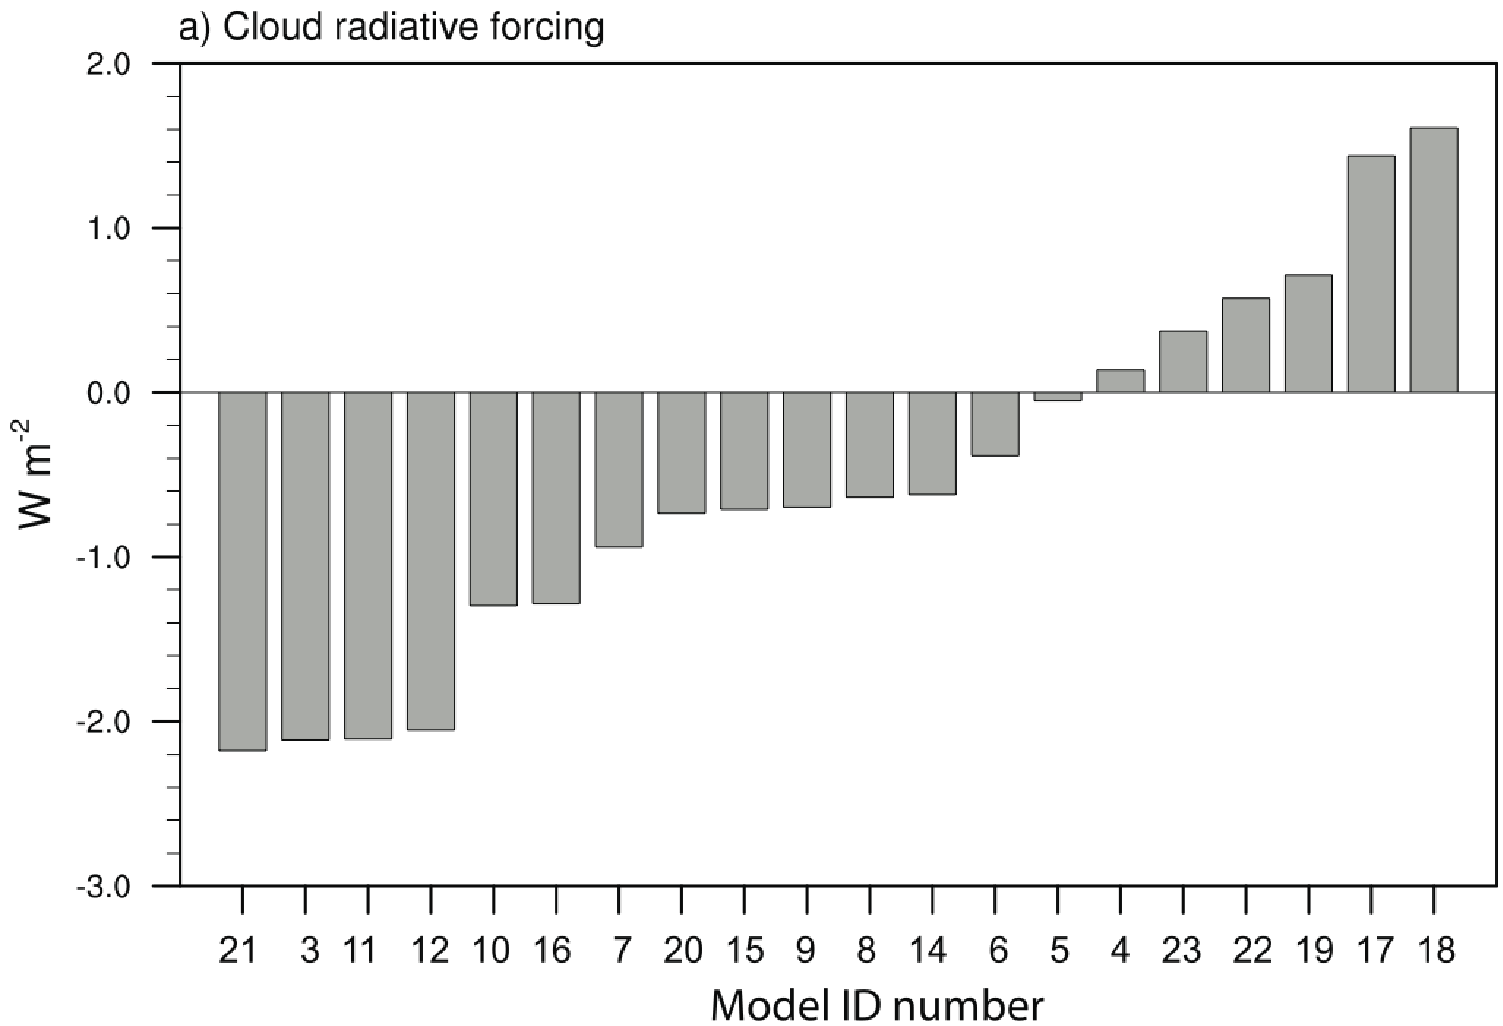

Regional Climate Change

5. On the regional level, uncertainties in regional climate change, e.g. whether we can expect more floods or more droughts, more storms or longer periods of unsettled weather, are even bigger.

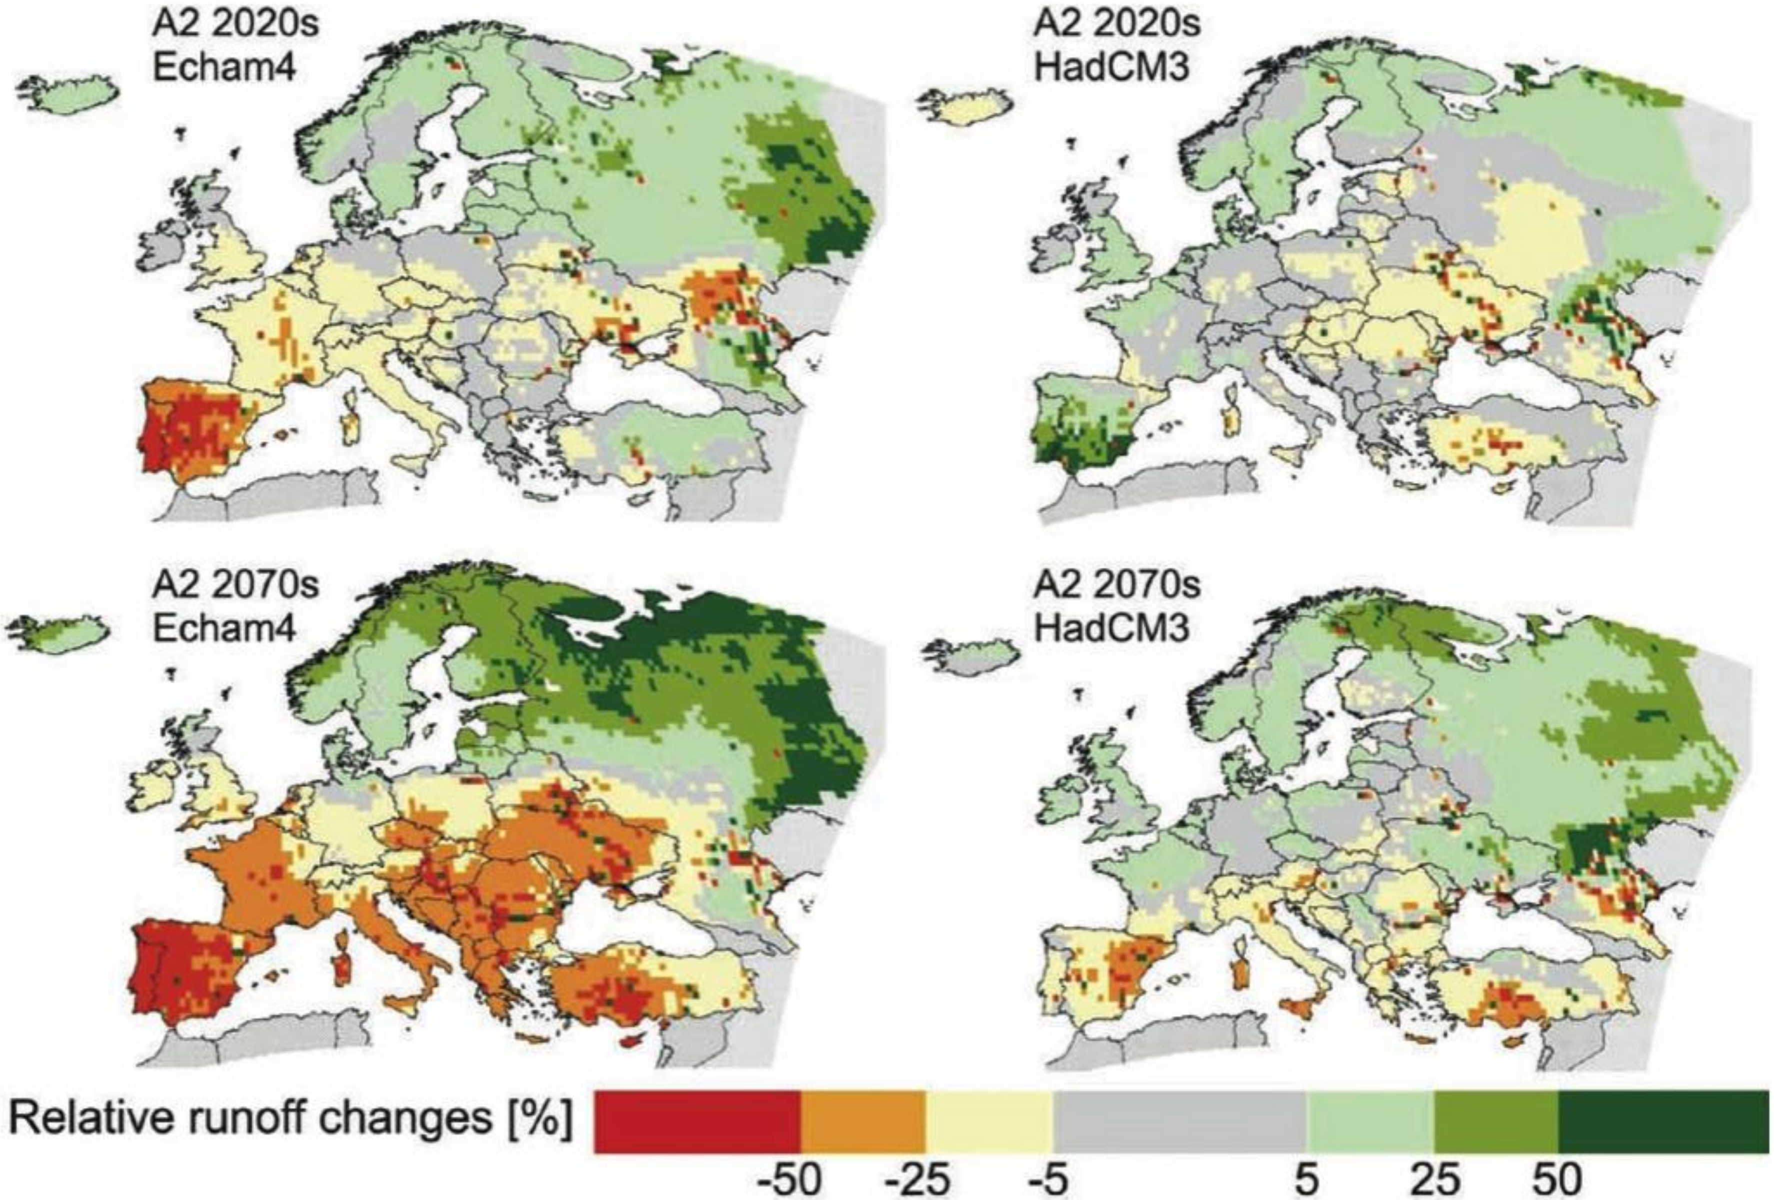

Figure 12.1. *Change in annual river runoff between the 1961-1990 baseline period and two future time slices (2020s and 2070s) for the A2 scenarios (Alcamo et al., 2007).*

What will be the time-evolving regional/local climate changes (from now until 2099 and beyond) to which human societies will have to adapt?

$\Rightarrow$ Post-AR4 climate change modeling paradigm:

- **Decadal Prediction**
- Better regional predictions of weather and climate extremes
- Higher resolution, initialized AOGCMs
- To advise adaptation decisions

- **Longer Term Projections**
- New mitigation scenarios to address processes and feedbacks, e.g. carbon cycle
- Earth system models with intermediate resolution
- To advise mitigation and adaptation decisions

Climate Prediction - Computing

6. Climate prediction is one of the most computationally demanding problems.

Weather - Climate Continuum

- Instead of predicting weather 15 days in advance, need to predict climate 10 years in advance
- Initial value problem at all lead-time scales
  - Need to initialize entire physical climate system (atmosphere, ocean, ice, land surface)
  - Need to assimilate Earth observations
- Multi-scale, multi-physics problem at all lead-time scales

Longer Runs

Longer Runs

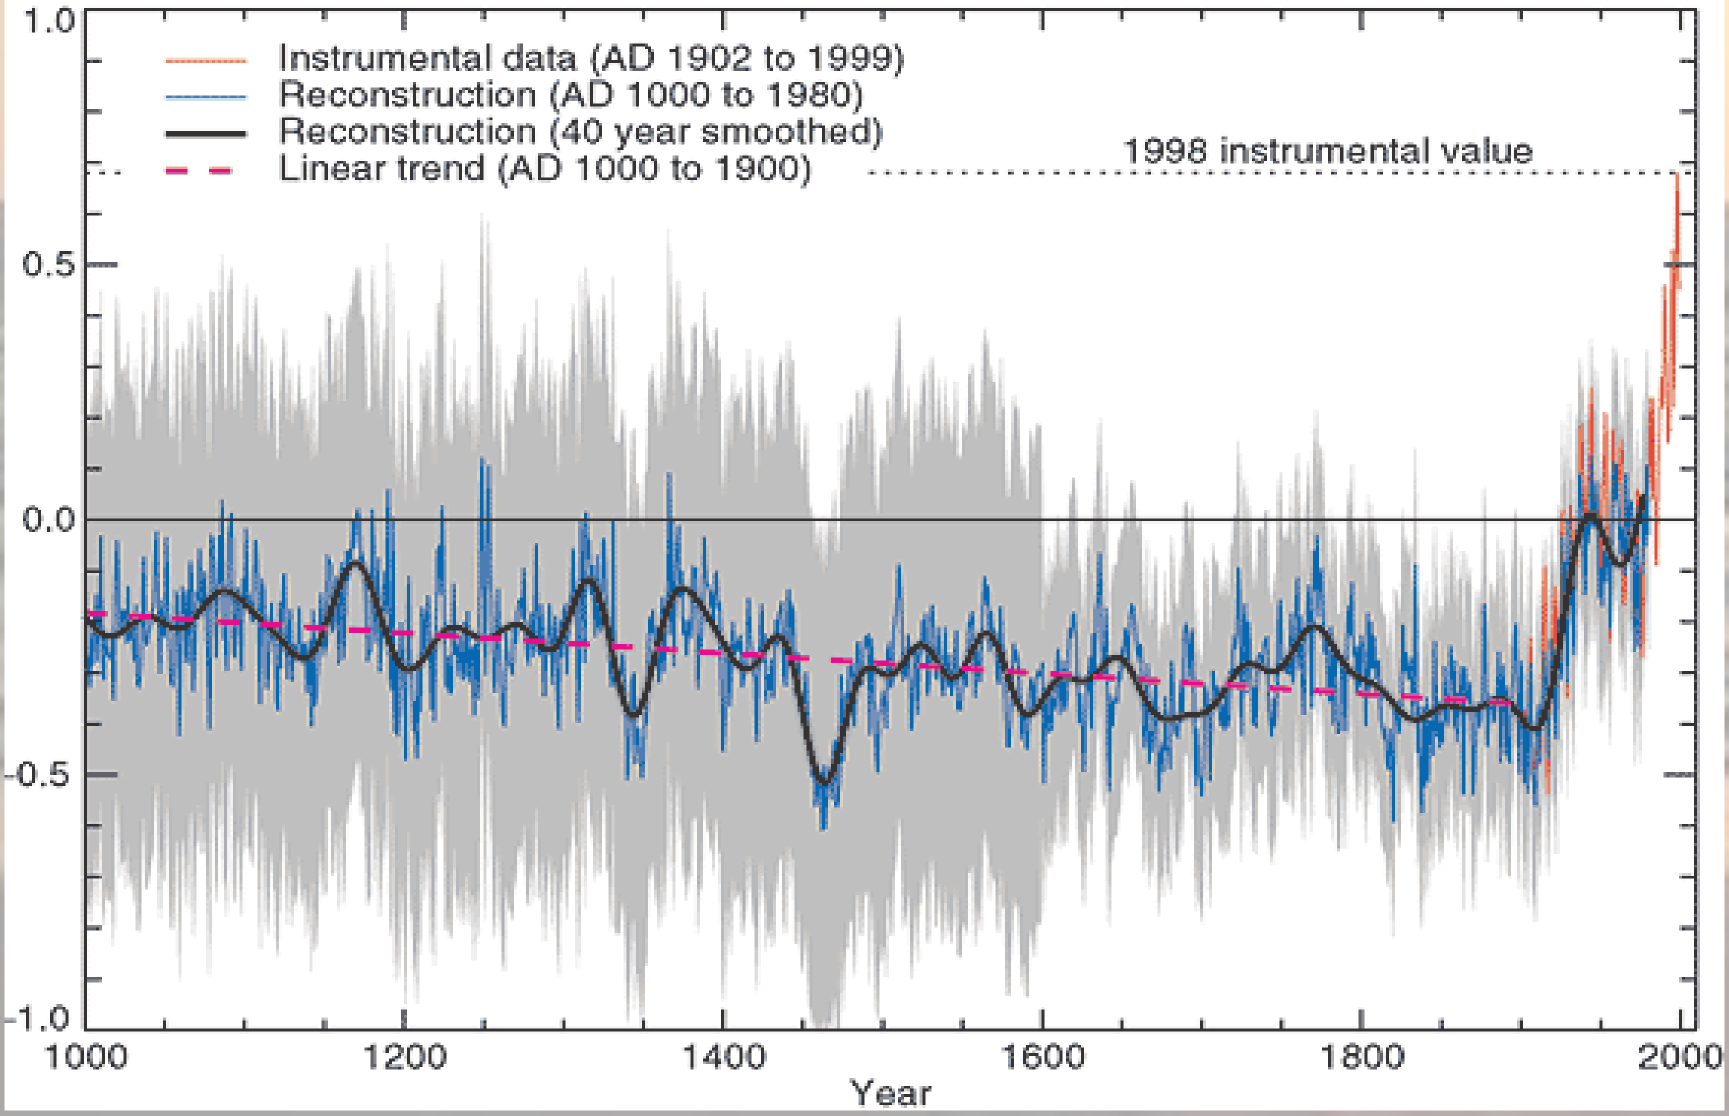

Instrumental data (AD 1902 to 1999)
Reconstruction (AD 1000 to 1980)
Reconstruction (40 year smoothed)
Linear trend (AD 1000 to 1900)
1998 instrumental value
1.0
0.5
0.0
-0.5
-1.0
1000
1200
1400
1600
1800
2000
Year

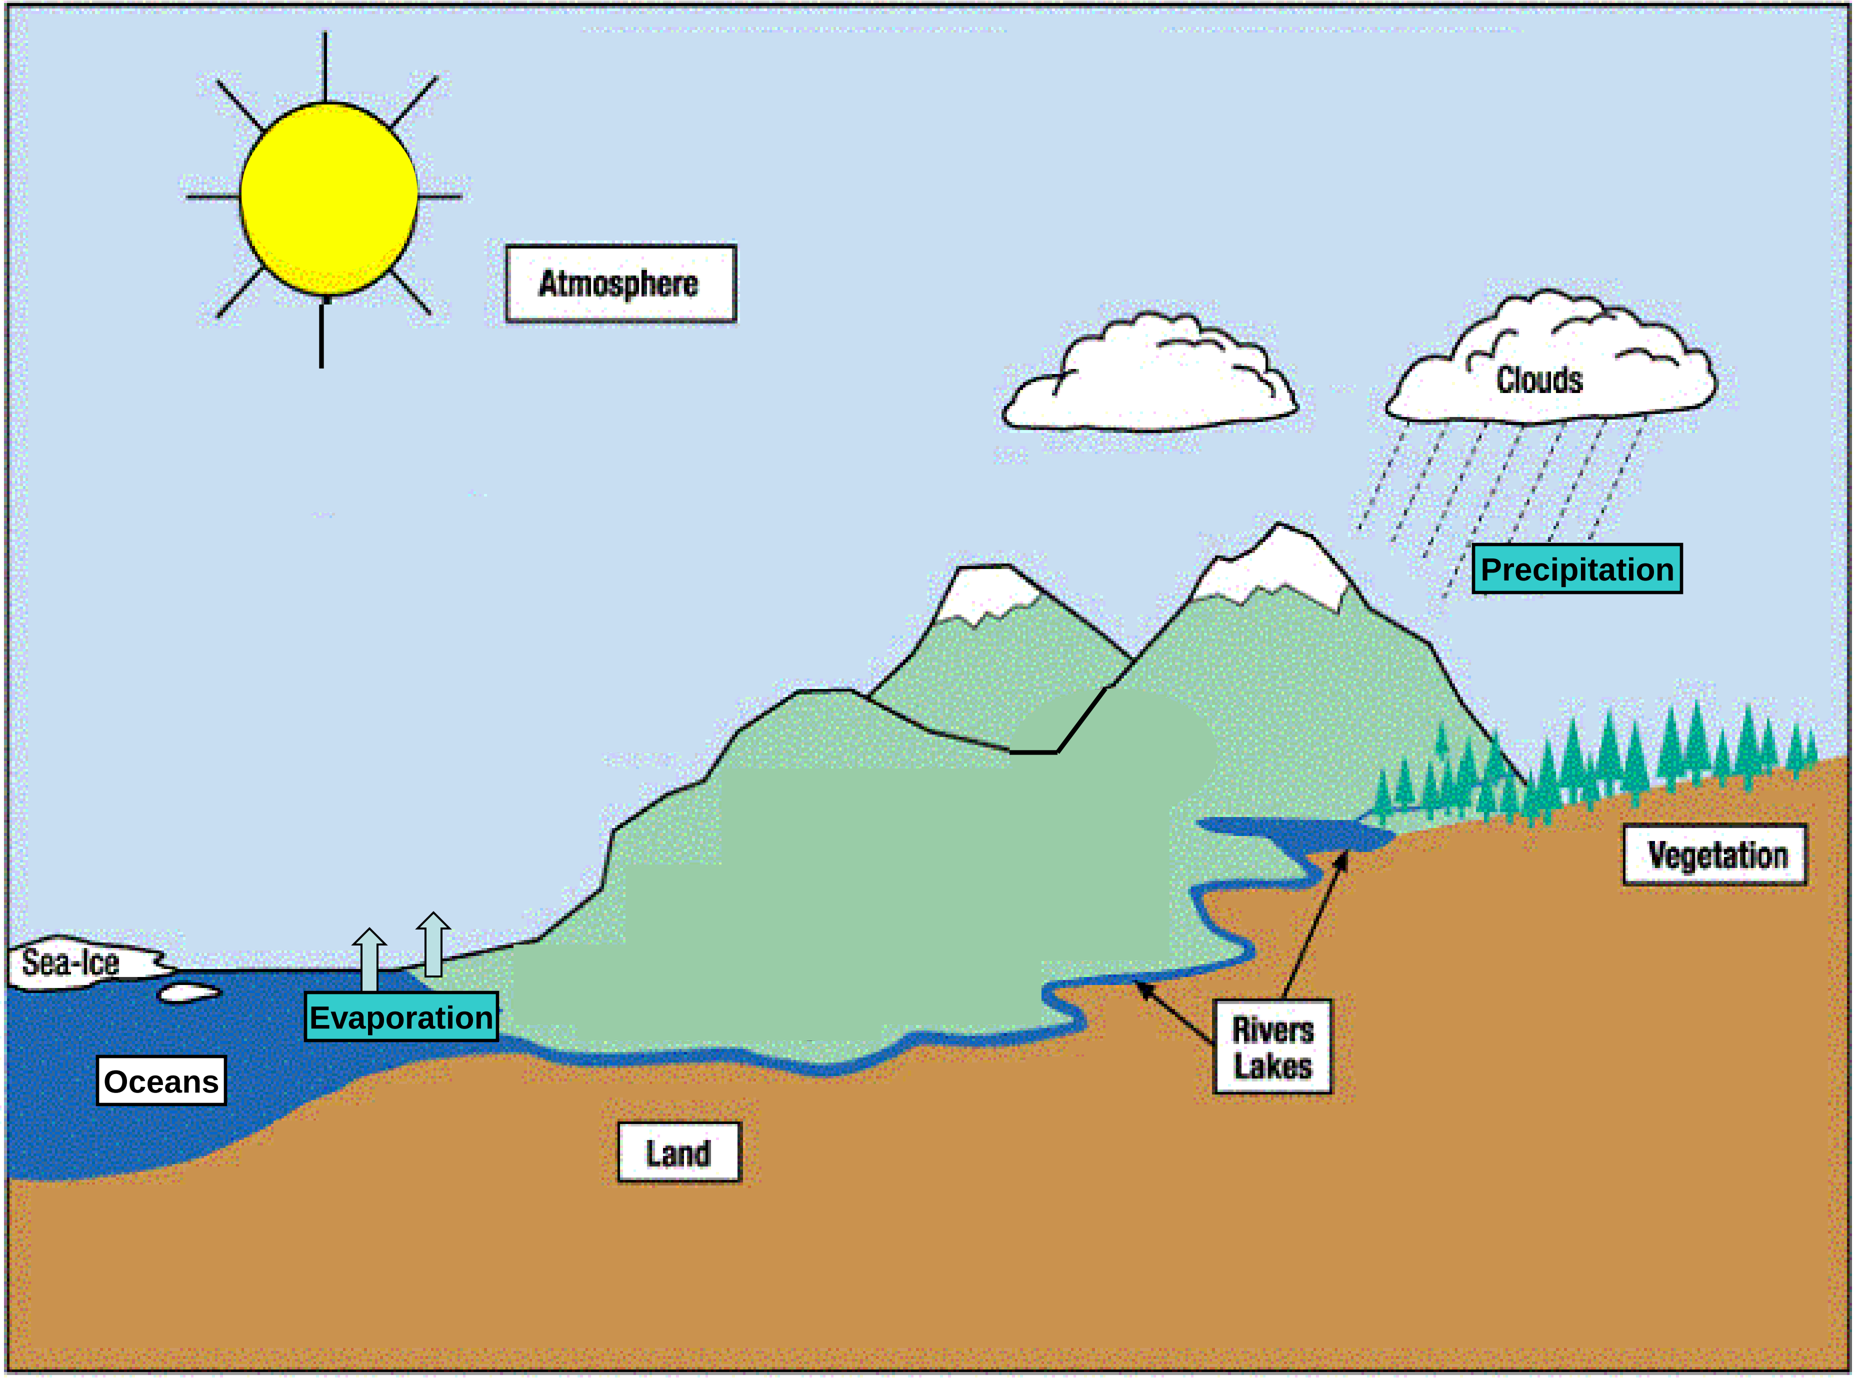
Atmosphere
Clouds
Precipitation
Sea-Ice
Evaporation
Oceans
Land
Rivers
Lakes
Vegetation

Feedbacks are essential

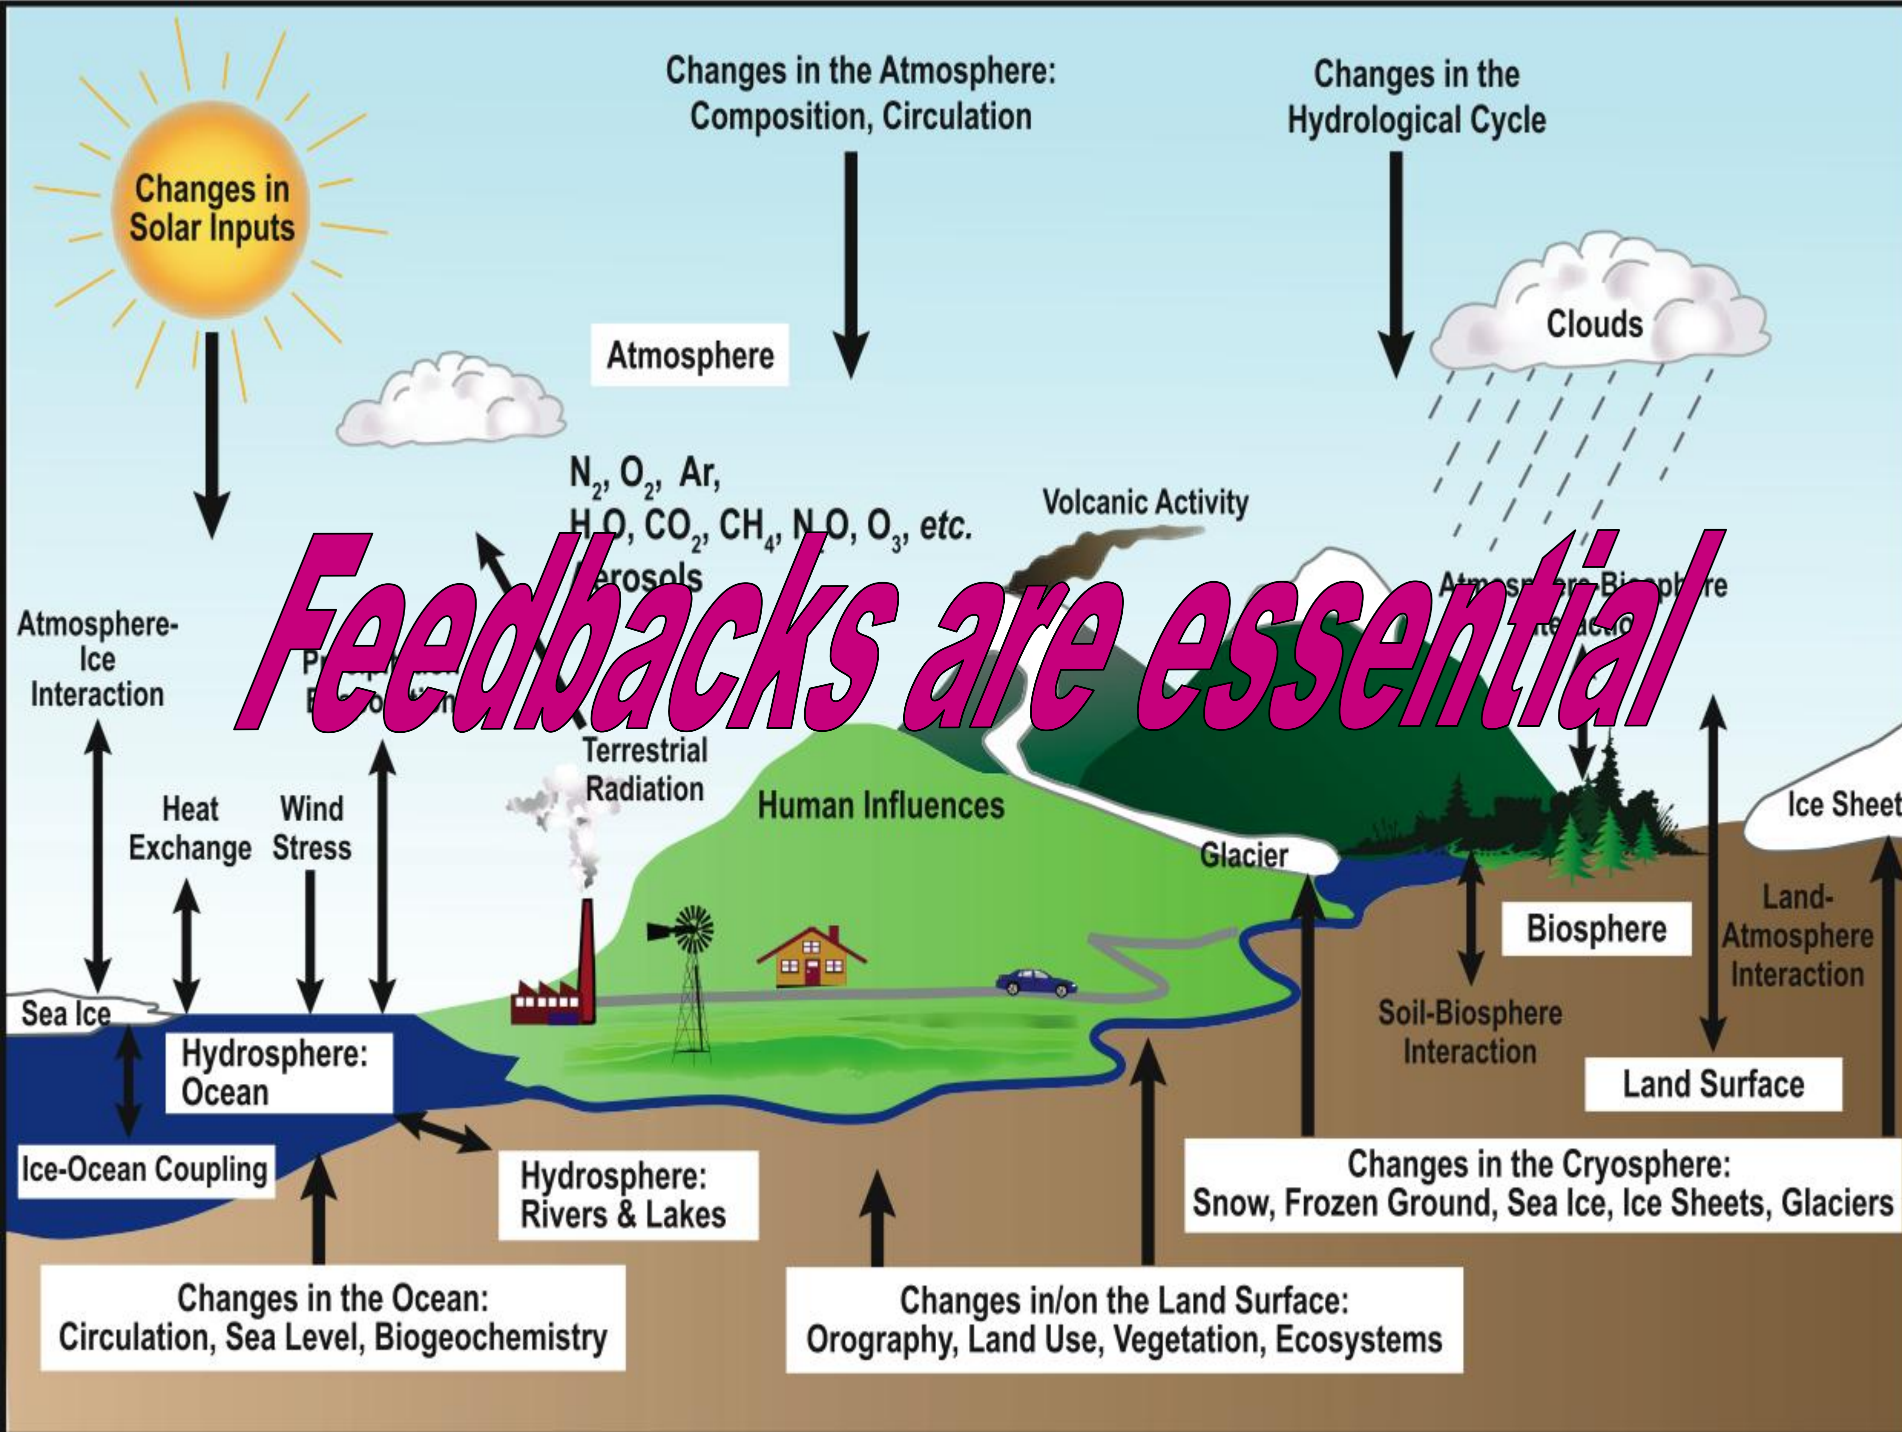

Ensembles

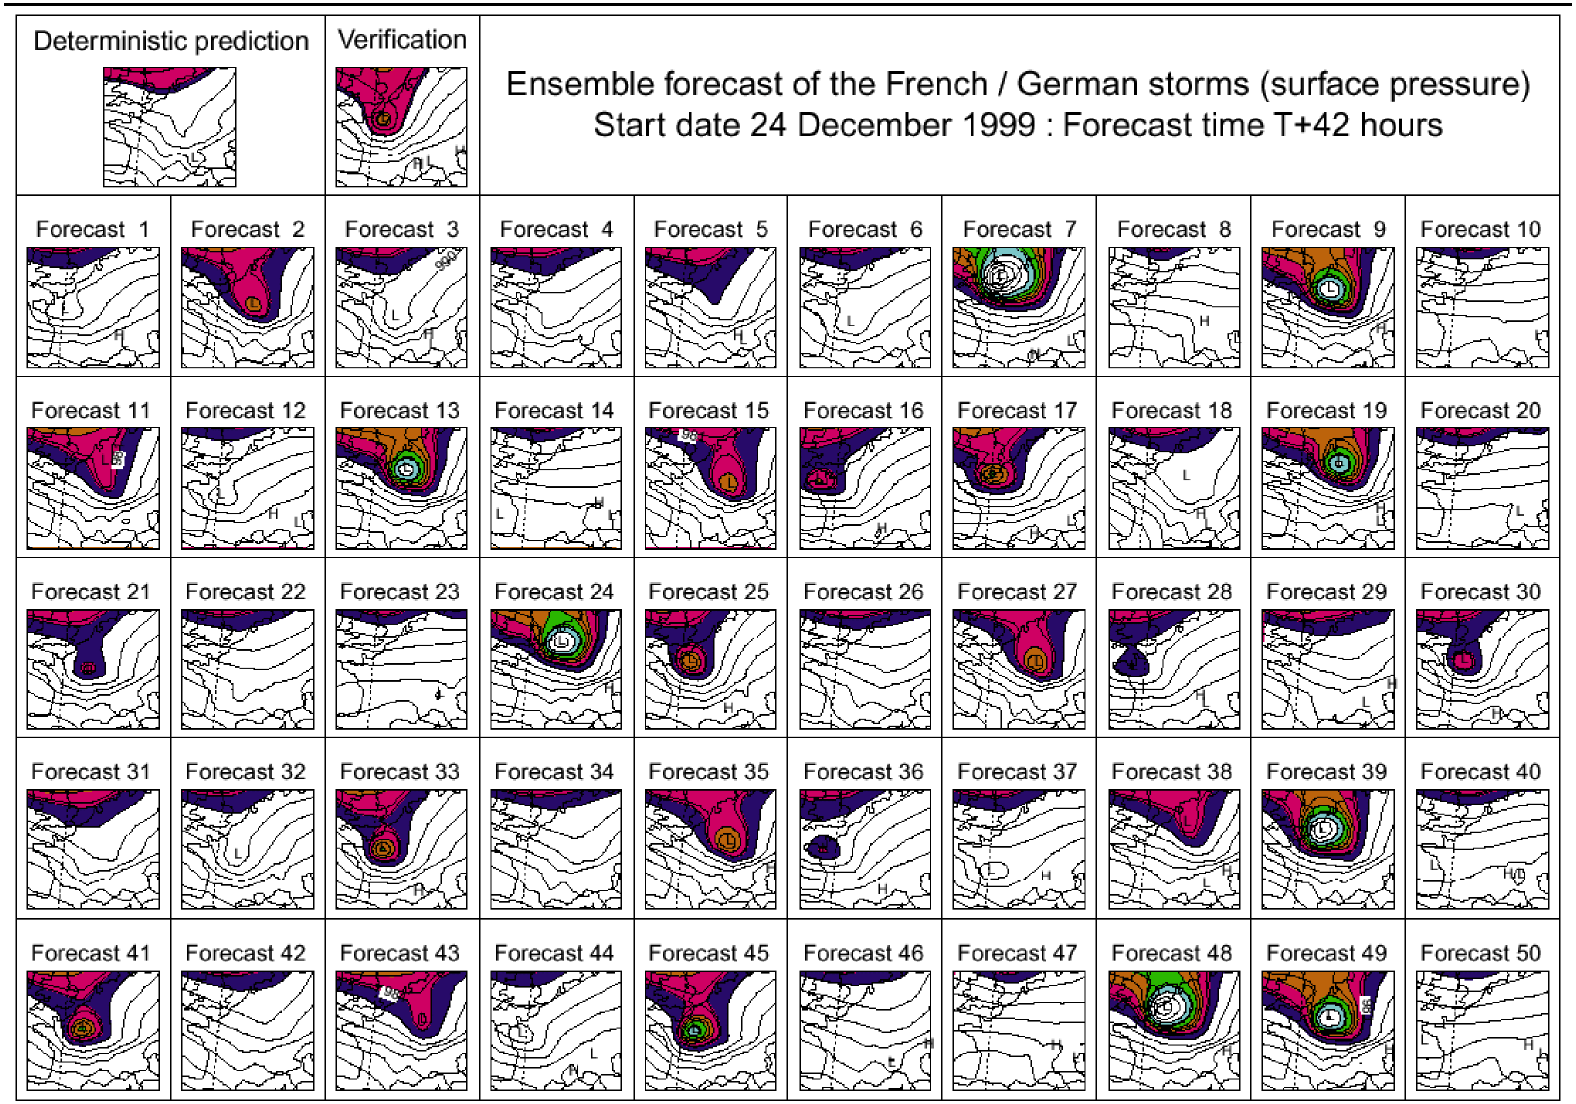

Courtesy of T. Palmer (ECMWF)

Multi-Model Ensembles

CMIP3 archive at PCMDI:

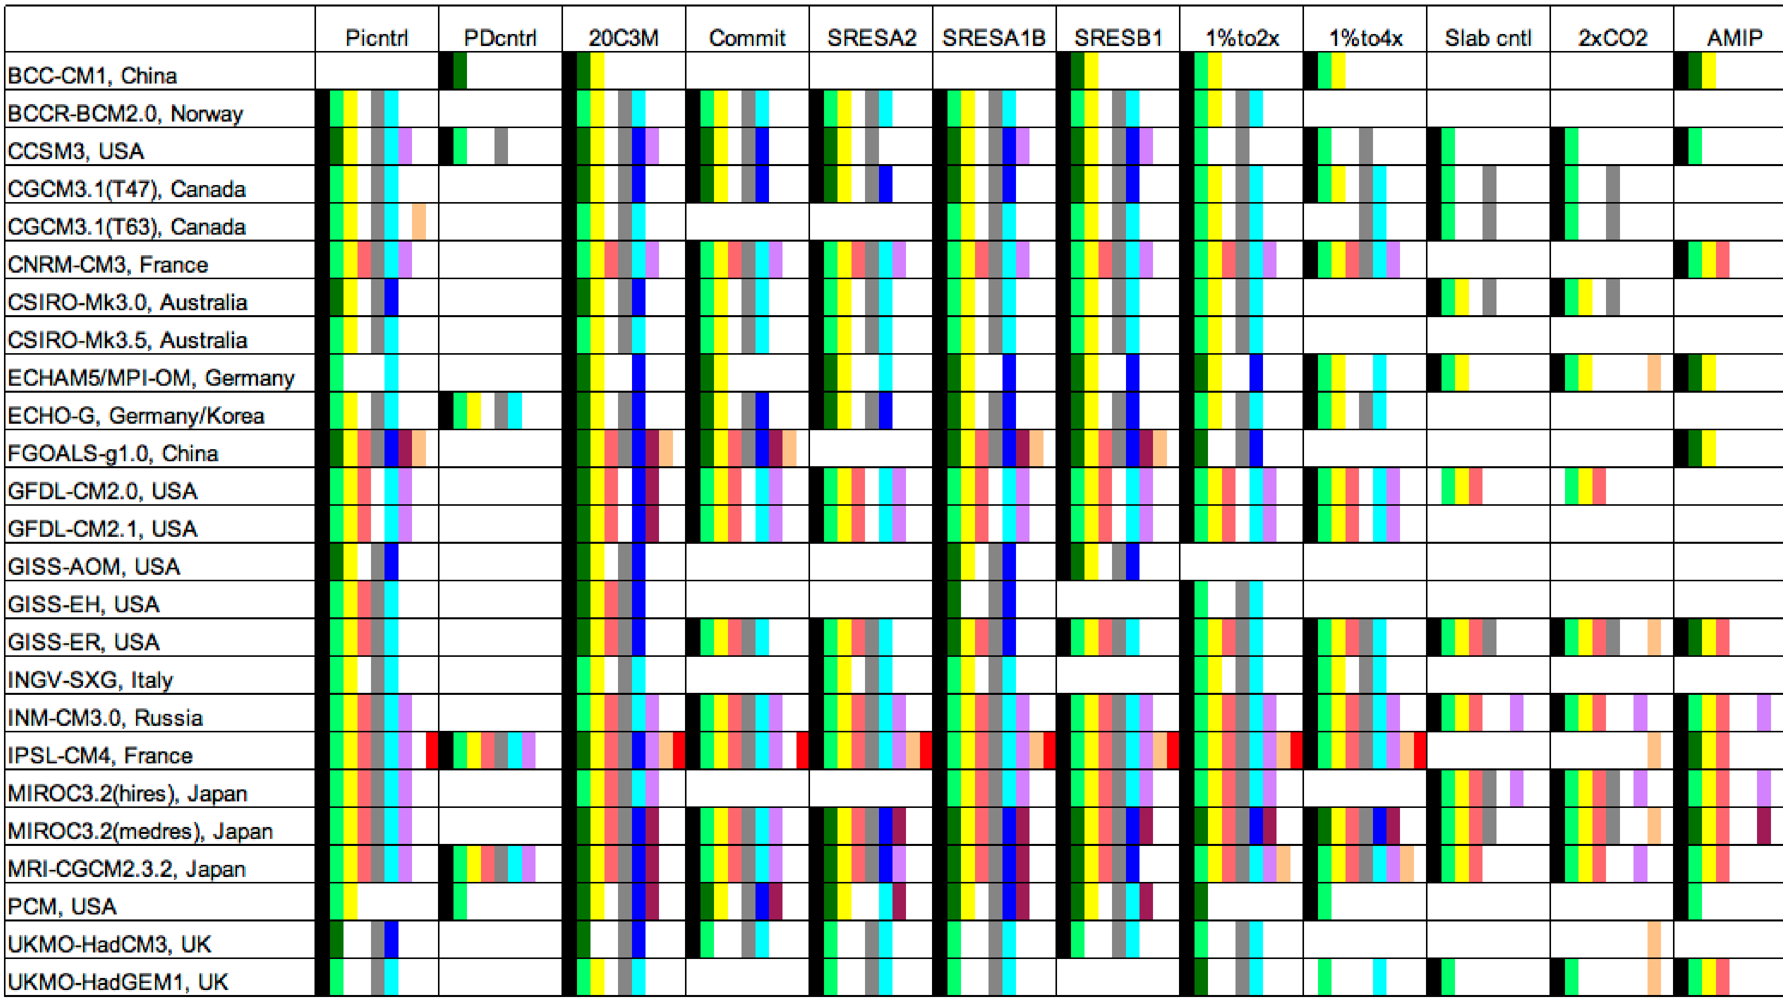

Spatial Resolution

7. The spatial scales of physical processes important to climate span almost 10 decades; from the global (10^7 m) to the viscous (10^{-2} m), but even the largest climate models resolve less than 4 decades.

Spatial Resolution…

Spatial Resolution…

- ➲ is required for:
  - Accuracy
  - Representation of features
  - Representation of processes
  - Scale interaction
- ➲ drives computational demand

Topography (m; res = 20km)

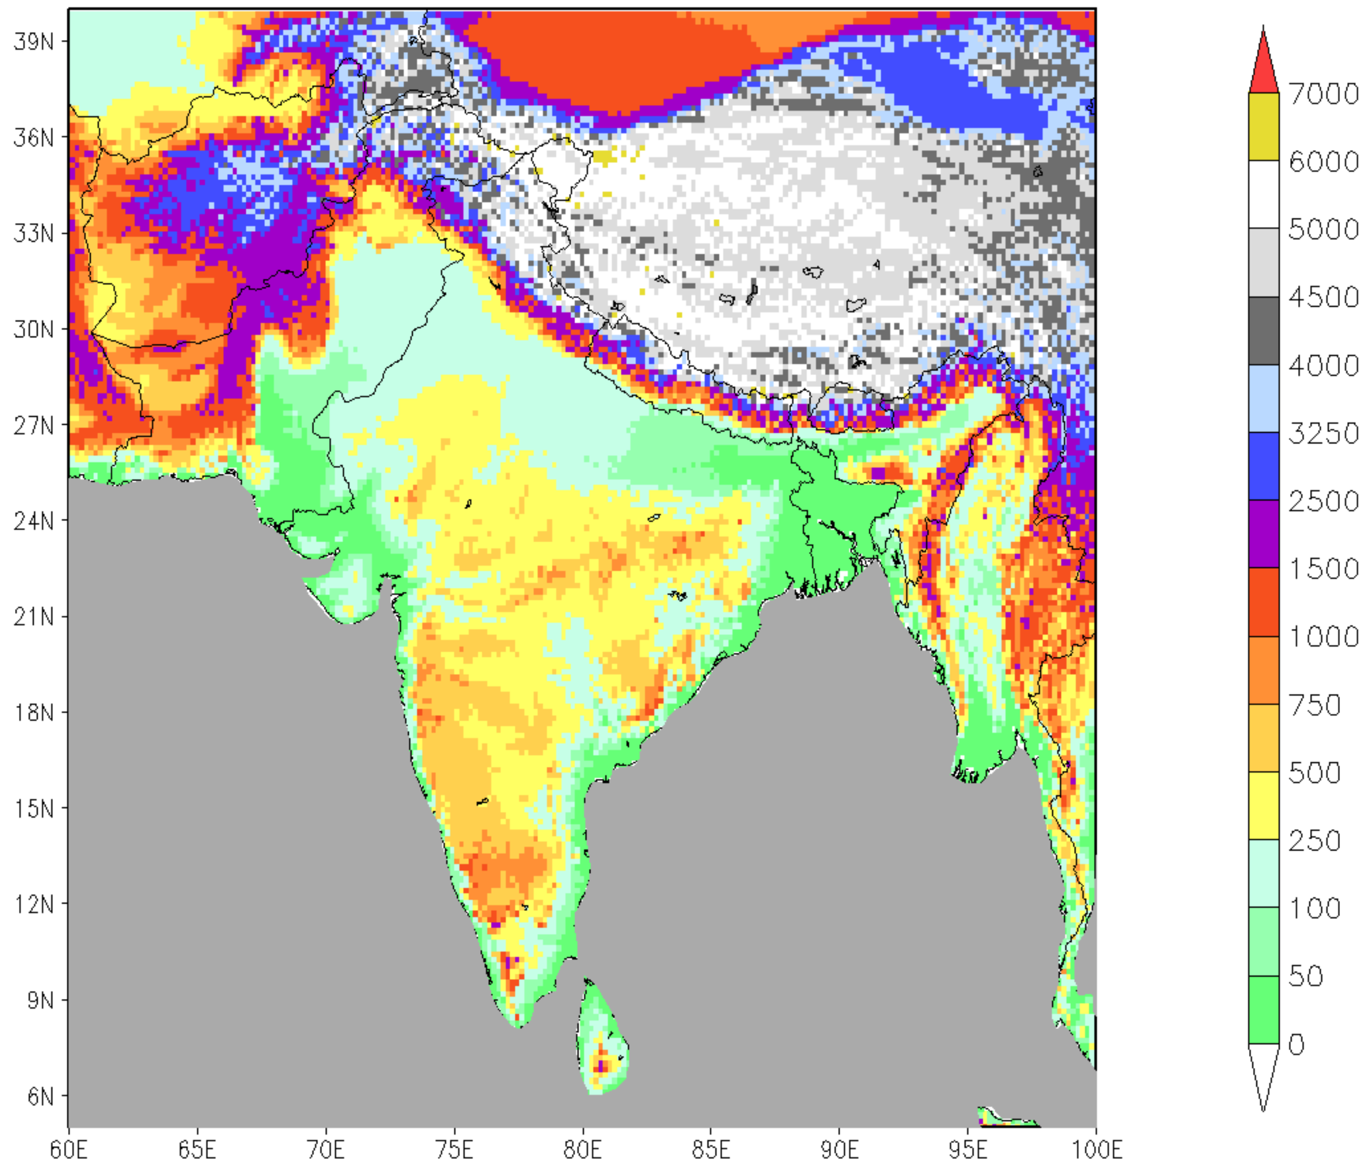

CCSM
CAM @ 0.5º + POP @ 1.0º

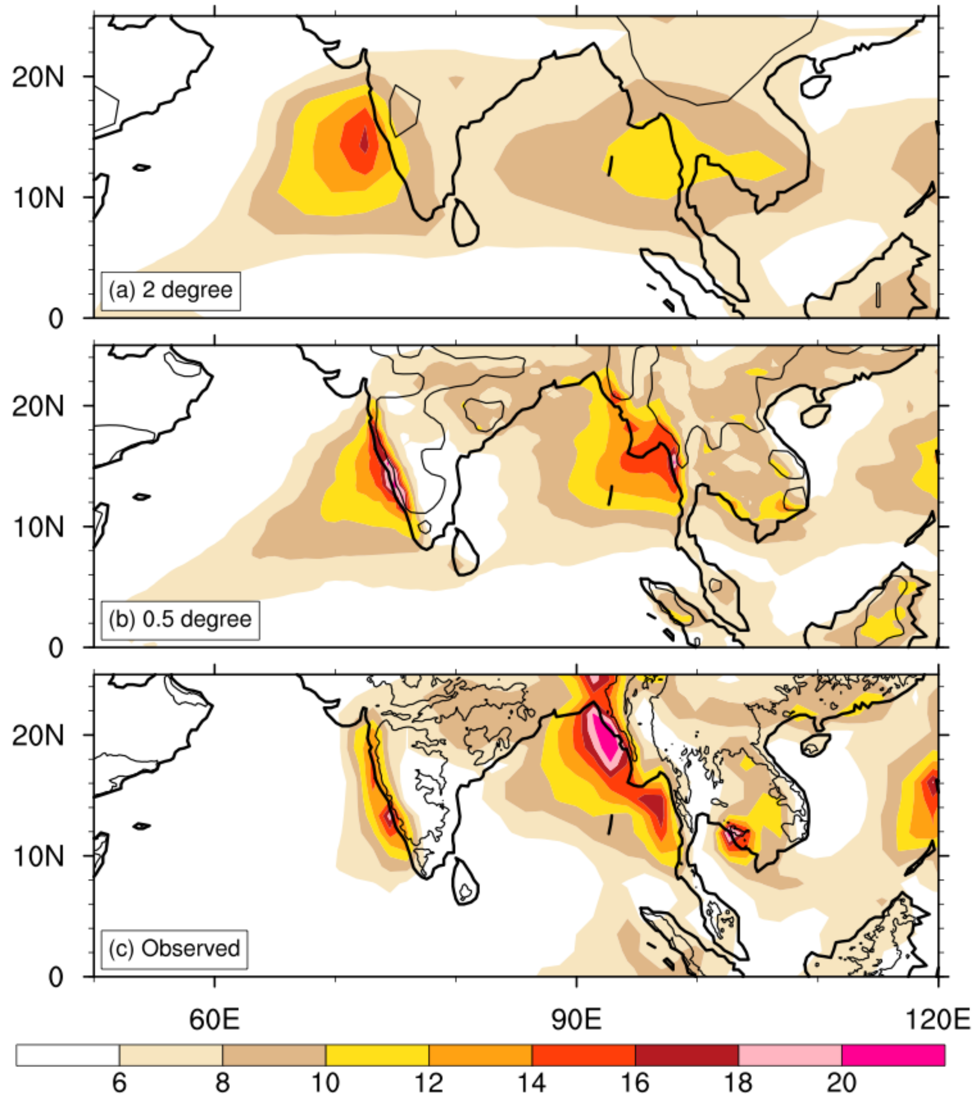

Gent et al. 2009 (*Clim. Dyn.*)

Climate Problem: Predicting Tropical Cyclone Characteristics

- Frequency
- Regions affected (tracks)
- Intensity and areal extent (wind, waves, surge)
- Tornadoes!!
- Requires:
  - High model resolution
  - Improved physics
    - Cloud processes
    - Air-sea exchange
  - Combination of model and statistics

Courtesy of Greg Holland - Sep 2007

Hurricane Katrina Intensity at Landfall

29 Aug 2005 14 Z

4 km WRF, 62 h forecast

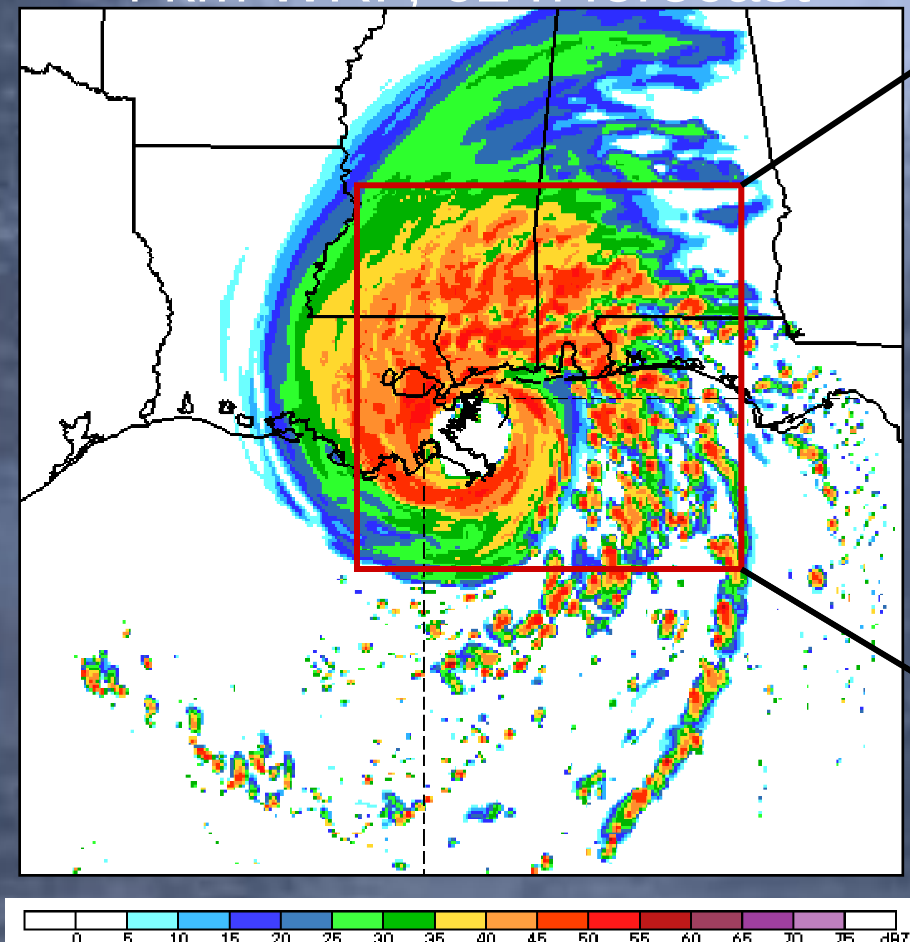

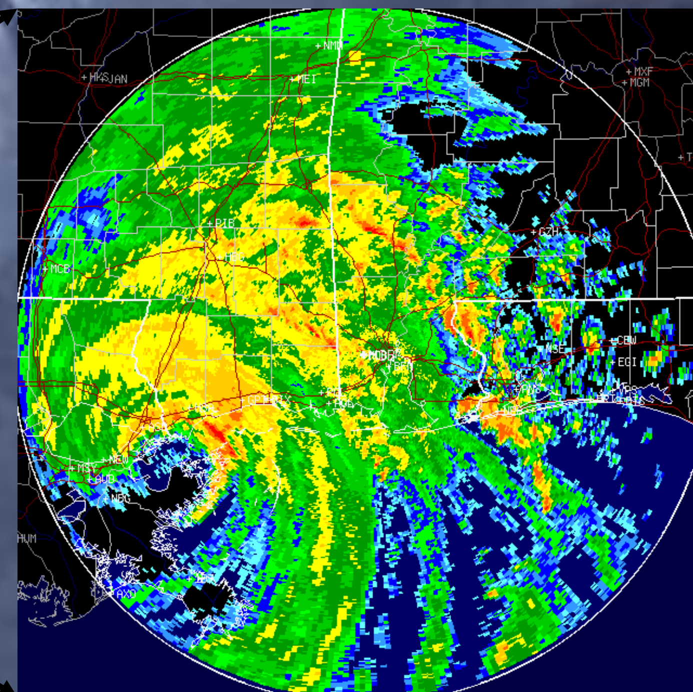

Mobile Radar

Courtesy of P. Fox (NCAR)


Need Long Runs!

Accuracy / Resolution Issues in the Atmosphere

- **Convective parameterization - when does it help? hurt?**
  - **Convective parameterization tends to reduce strength of tropical storms, cloud time scale is key parameter; may be better off without CU parameterization (GMAO finding)**
  - **35-km resolution needed to get seasonal cycle of tropical storm frequency (NCEP finding)**
  - **Crossover resolution at which no CU parameterization is better may be ~ 20 km (GFDL hypothesis)**
  - **Evidence that cloud-resolving models are qualitatively different …**
- **50-km or even 25-km resolution is needed to get statistics of extratropical winter storms right (ECMWF finding)**
- **17-km resolution may be required to get important extreme events in extratropics right (ECMWF finding)**


Need Process-Resolving Models!

Global Cloud-System Resolving Model (NICAM)

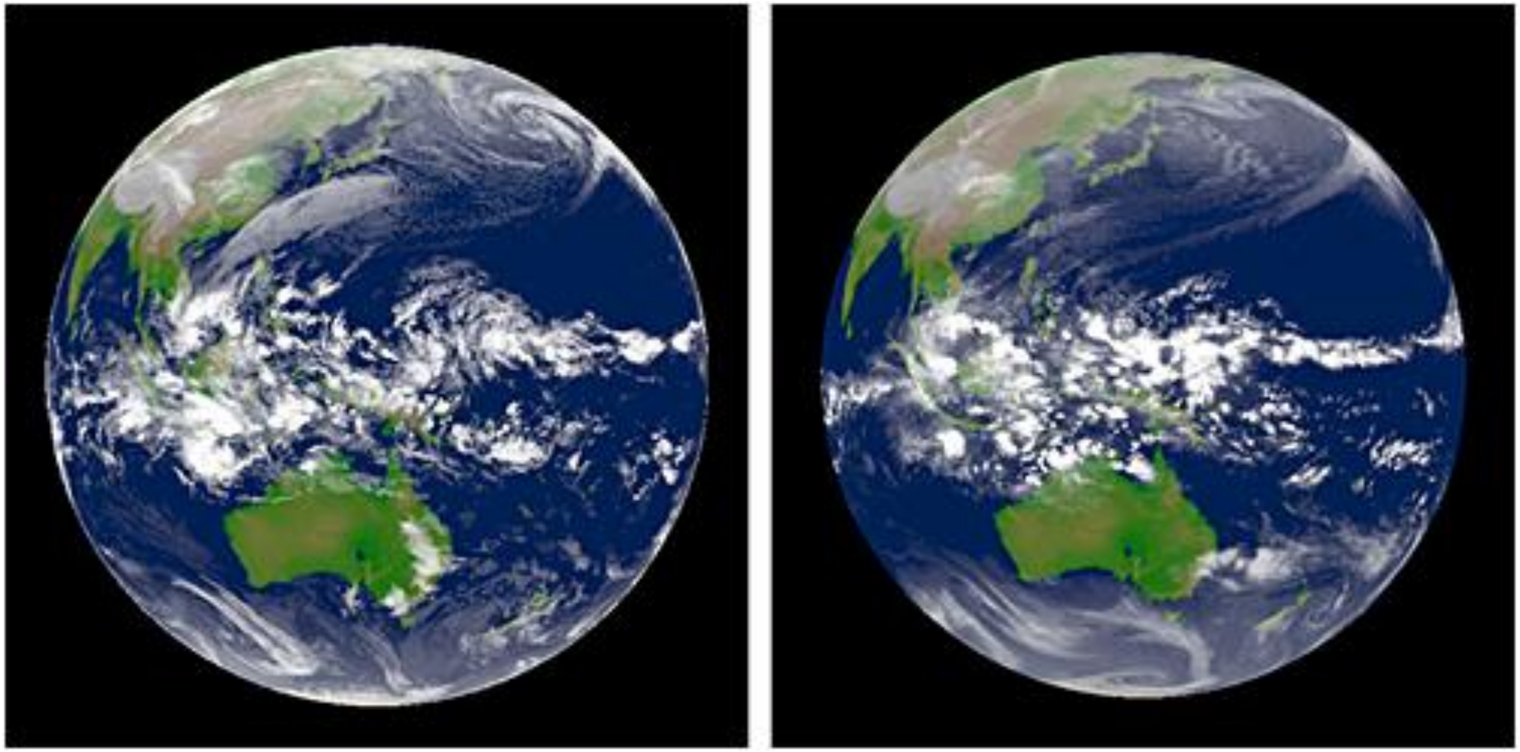

Snapshot of 00UTC31Dec2006 cloud distribution in satellite image (left) and 3.5-km NICAM simulation (right)

Miura et al., 2007

MJO Simulation

(Winter, 20-100 bandpass)

SP-CCSM **GPCP/NCEP** **CCSM**

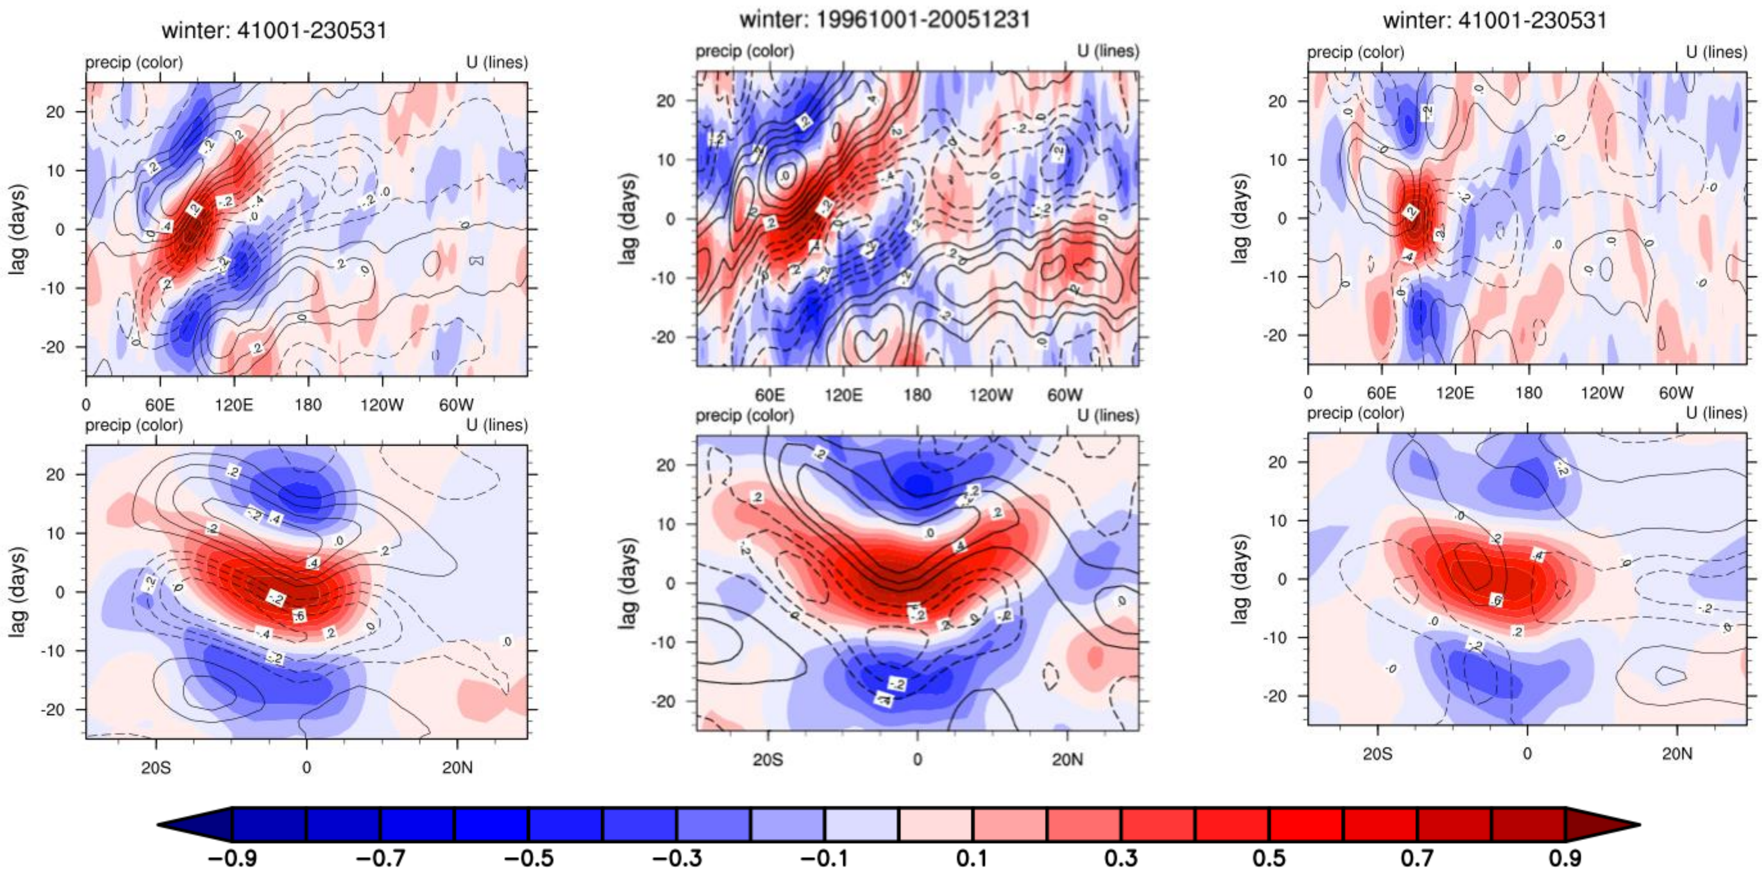

Courtesy Cristiana Stan, COLA

Accuracy / Resolution Issues in the Atmosphere

- Convective parameterization - when does it help? hurt?
  - Convective parameterization tends to reduce strength of tropical storms, cloud time scale is key parameter; may be better off without CU parameterization (GMAO finding)
  - 35-km resolution needed to get seasonal cycle of tropical storm frequency (NCEP finding)
  - Crossover resolution at which no CU parameterization is better may be ~ 20 km (GFDL hypothesis)
  - Evidence that cloud-resolving models are qualitatively different …
- **50-km or even 25-km resolution is needed to get statistics of extratropical winter storms right (ECMWF finding)**
- **17-km resolution may be required to get important extreme events in extratropics right (ECMWF finding)**

Does Resolution Help with ENSO Features?

- Mean Climate - Tropical mean SST, precipitation
- Equatorial Pacific SST annual cycle
- NINO3.4 - Power spectrum, amplitude, phase w.r.t. annual cycle, correlation with wind stress, correlation with global SST
- ENSO events - Duration, propagation of SST anomalies, evolution of tropical Pacific thermocline
- Linkage to Subseasonal Variability
- Teleconnections - ENSO-monsoon correlation, mid-latitude stationary waves

ENSO Simulation
NINO3.4 (5S-5N,170W-120W)

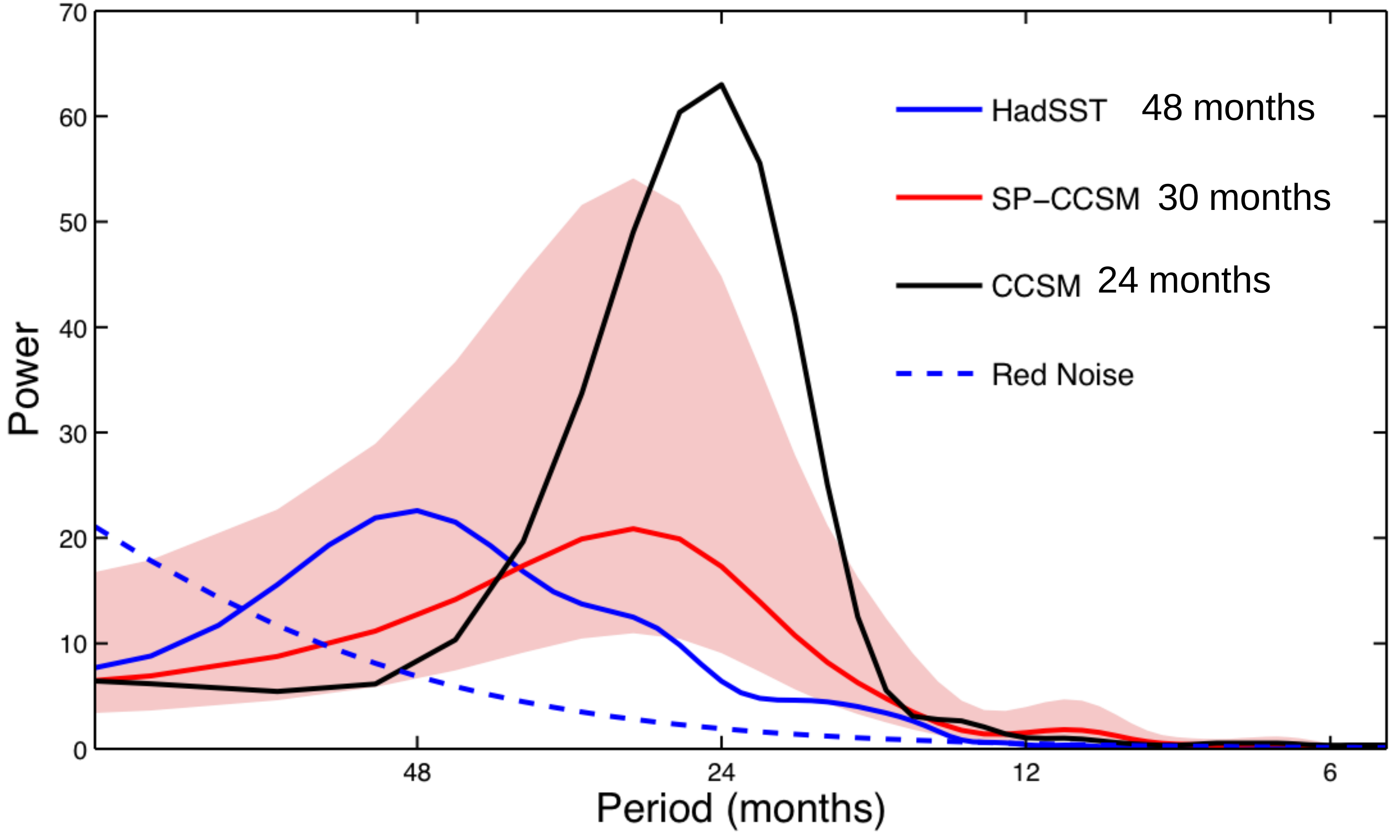

Courtesy Cristiana Stan, COLA

Global Atmospheric Models Horizontal Resolution

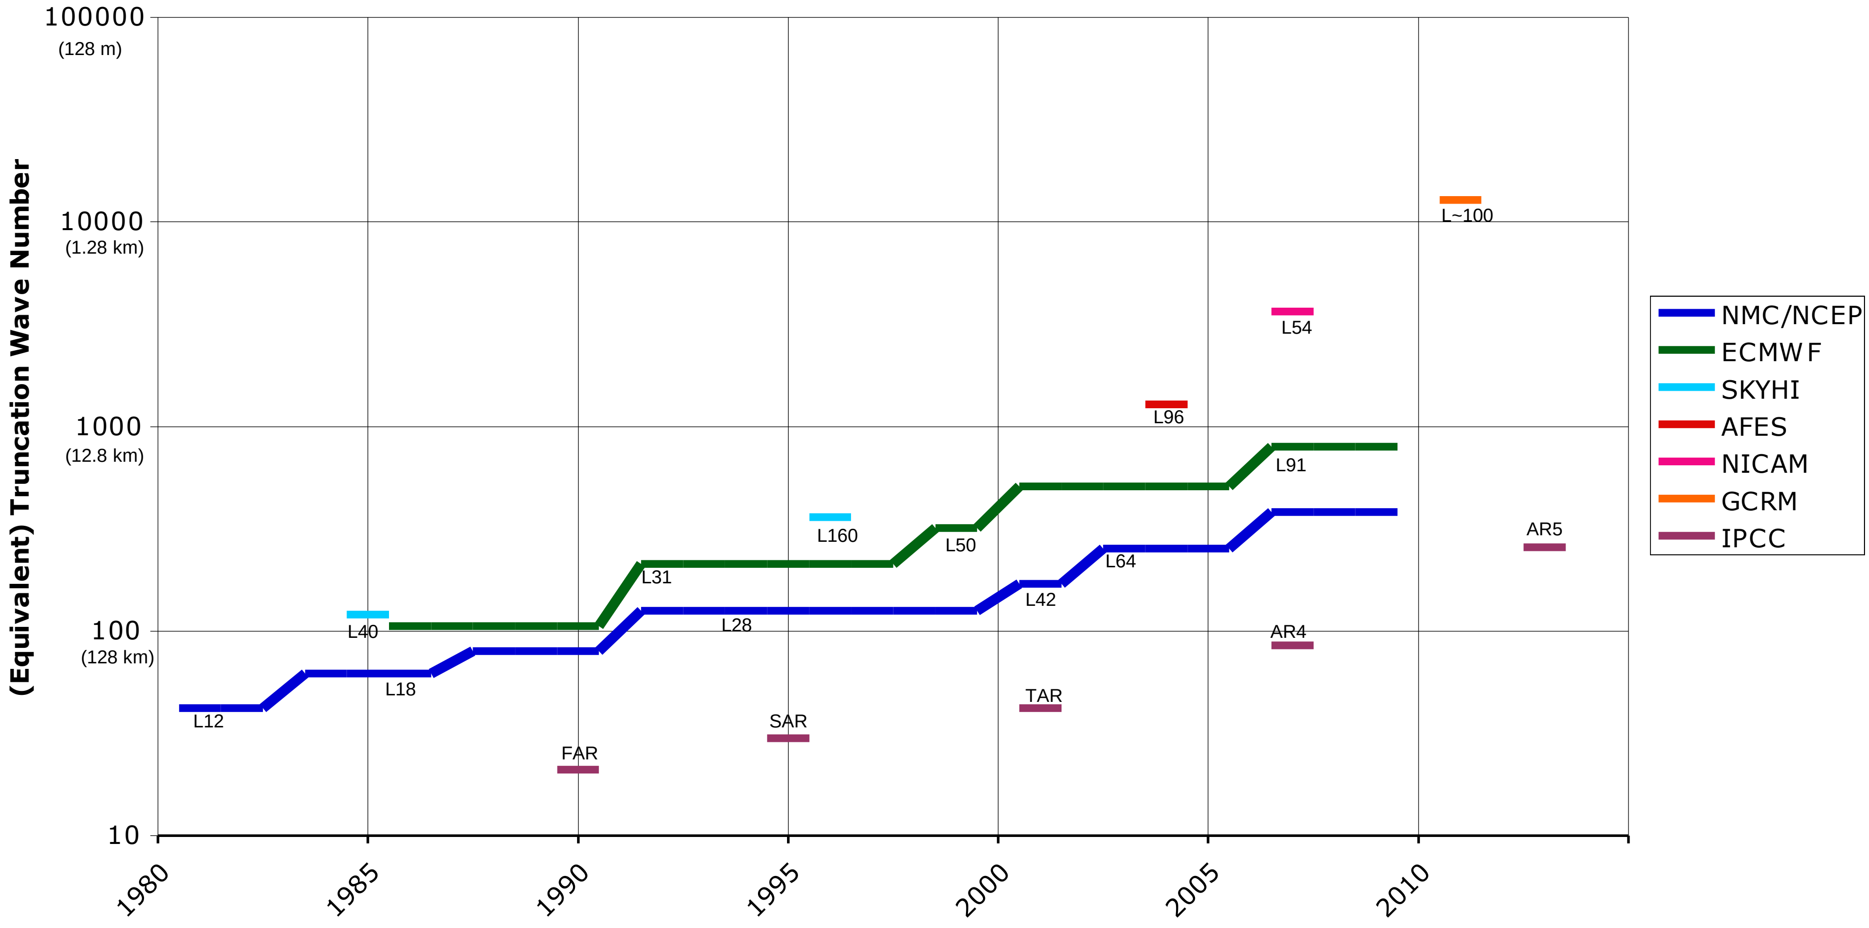

a la Kevin Hamilton (IPRC, 2009)

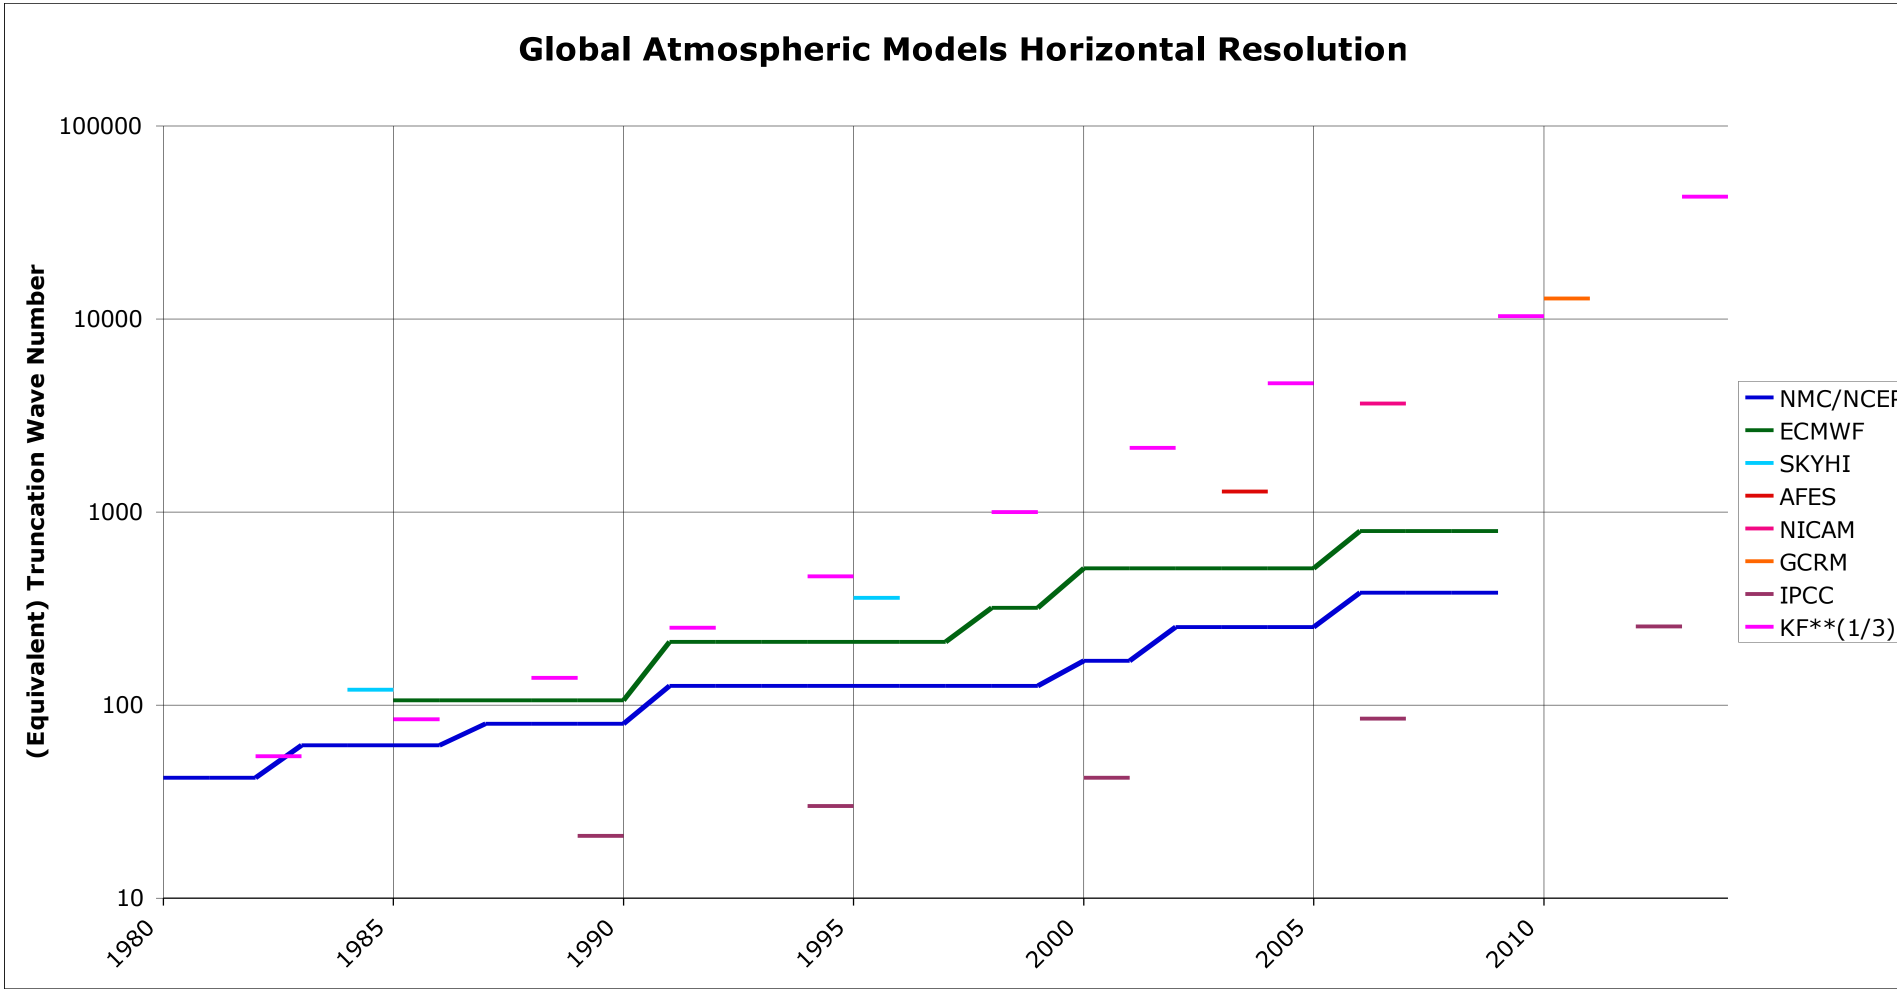

Global Atmospheric Models Horizontal Resolution
100000
10000
1000
100
10
(Equivalent) Truncation Wave Number
1980
1985
1990
1995
2000
2005
2010
NMC/NCEP
ECMWF
SKYHI
AFES
NICAM
GCRM
IPCC
KF**(1/3)

How Does Resolving the Ocean Circulation Help?

- **Typical global ocean models used for climate prediction “permit” eddies, at best**
- **To “resolve” eddies, need grid spacing less than 10 km (ocean modelers achieved this ~10 years ago)**

- **What is the impact of resolving the ocean eddies?**

… improves fidelity of ocean-only simulation:

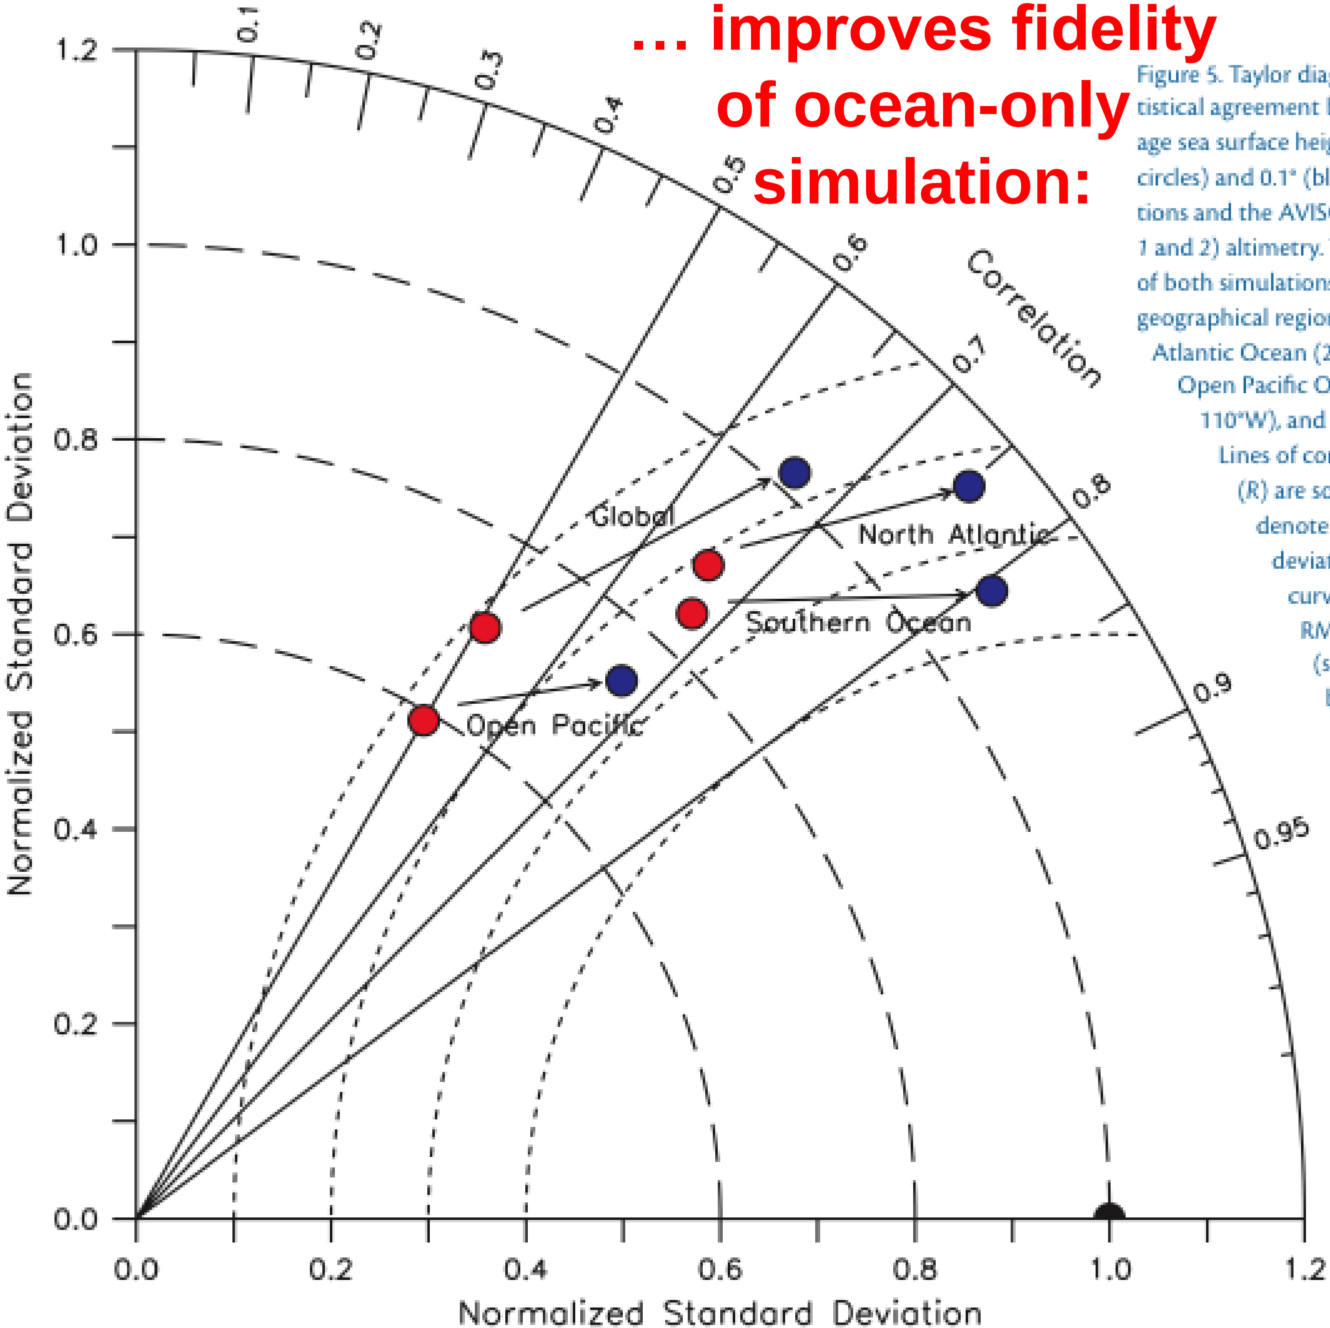

Figure 5. Taylor diagram showing the level of statistical agreement between the 1994–2001 average sea surface height anomaly from the 0.4° (red circles) and 0.1° (blue circles) global POP simulations and the AVISO (TOPEX/POSEIDON and *ERS* 1 and 2) altimetry. The arrows connect the results of both simulations evaluated over the following geographical regions: Global (70°S–70°N), North Atlantic Ocean (20°N–55°N, 100°W–20°W), Open Pacific Ocean (30°S–30°N, 150°E–110°W), and Southern Ocean (65°S–40°S). Lines of constant correlation coefficient (*R*) are solid; the long dashed curves denote lines of constant standard deviation ratio (σ); the short dashed curves denote lines of constant RMS difference, varying from 0.6 (small radius) to 0.9 by 0.1. The black semicircle represents the location of perfect agreement between the simulation and the comparison data set.

0.4º POP

0.1º POP

McClean, J. L., M. E. Maltrud, and F. O. Bryan, 2006: Measures of the fidelity of eddying ocean models. *Oceanography*, **19**, 104-117.

What About Coupled Models?

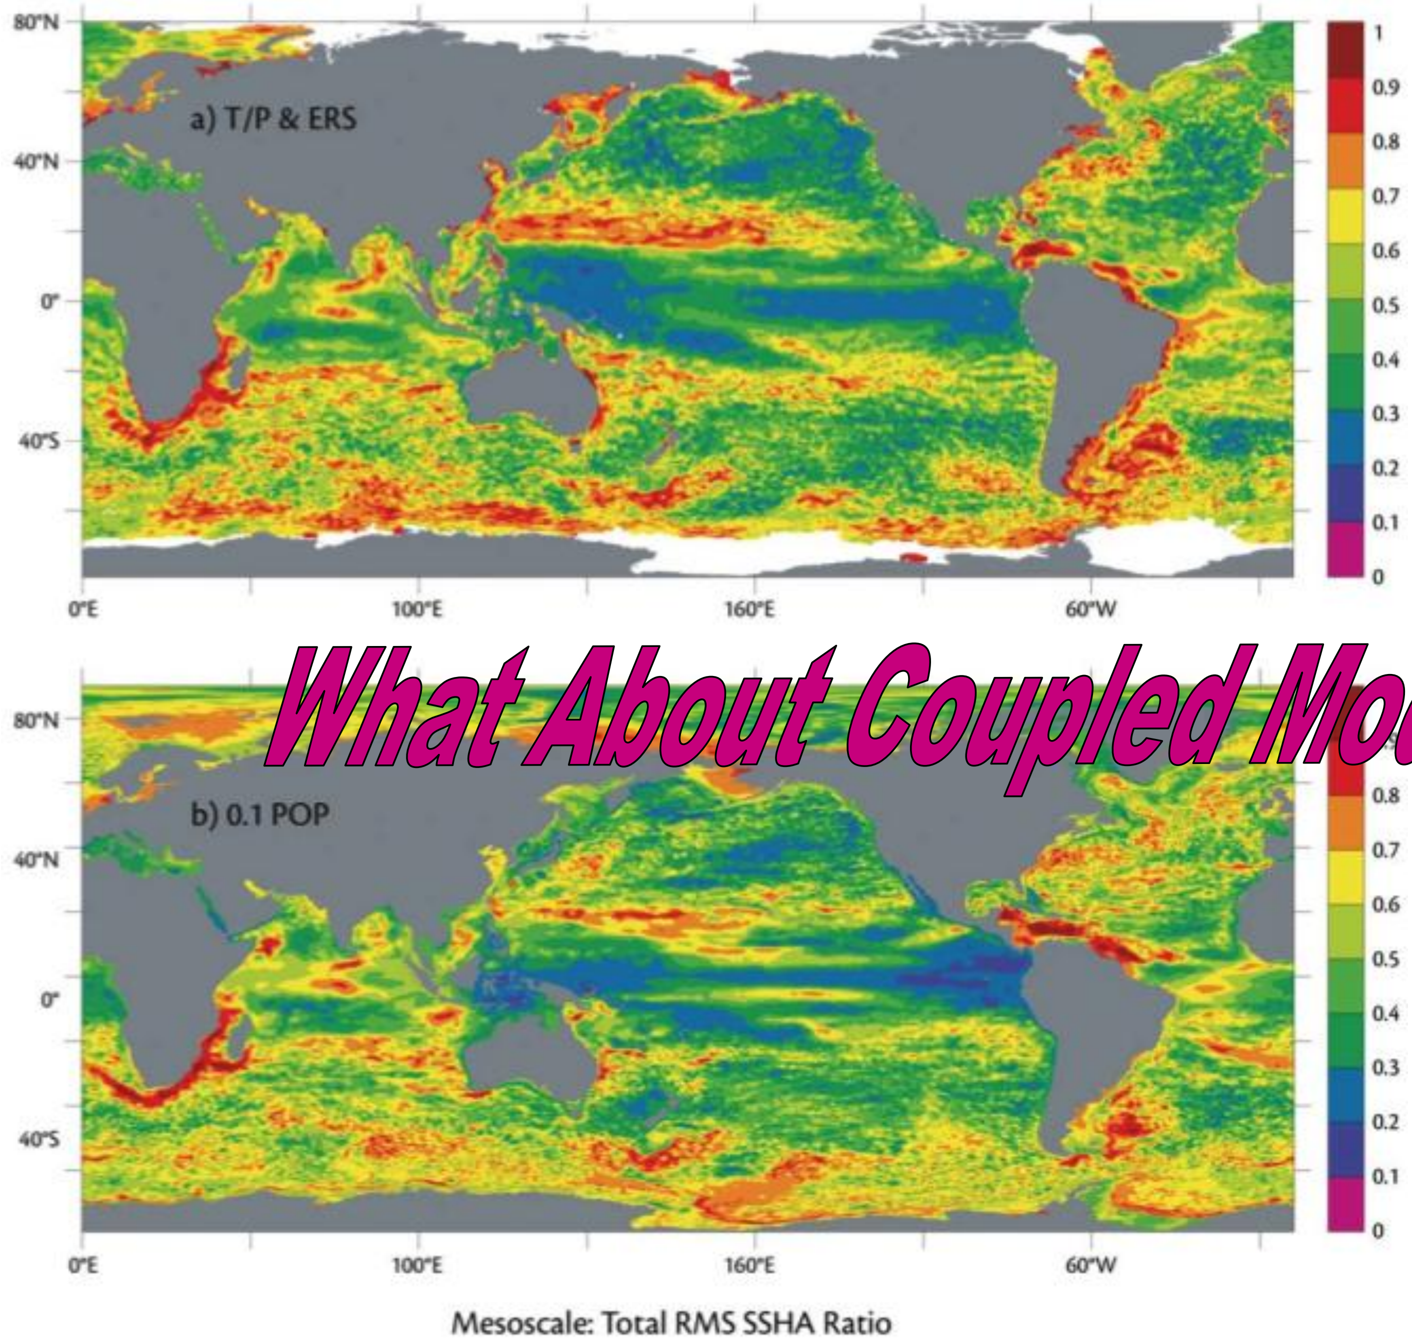

McClean, J. L., M. E. Maltrud, and F. O. Bryan, 2006: Measures of the fidelity of eddying ocean models. *Oceanography*, **19**, 104-117.

Correlation High-Pass SST vs. $|V_{srf}|$

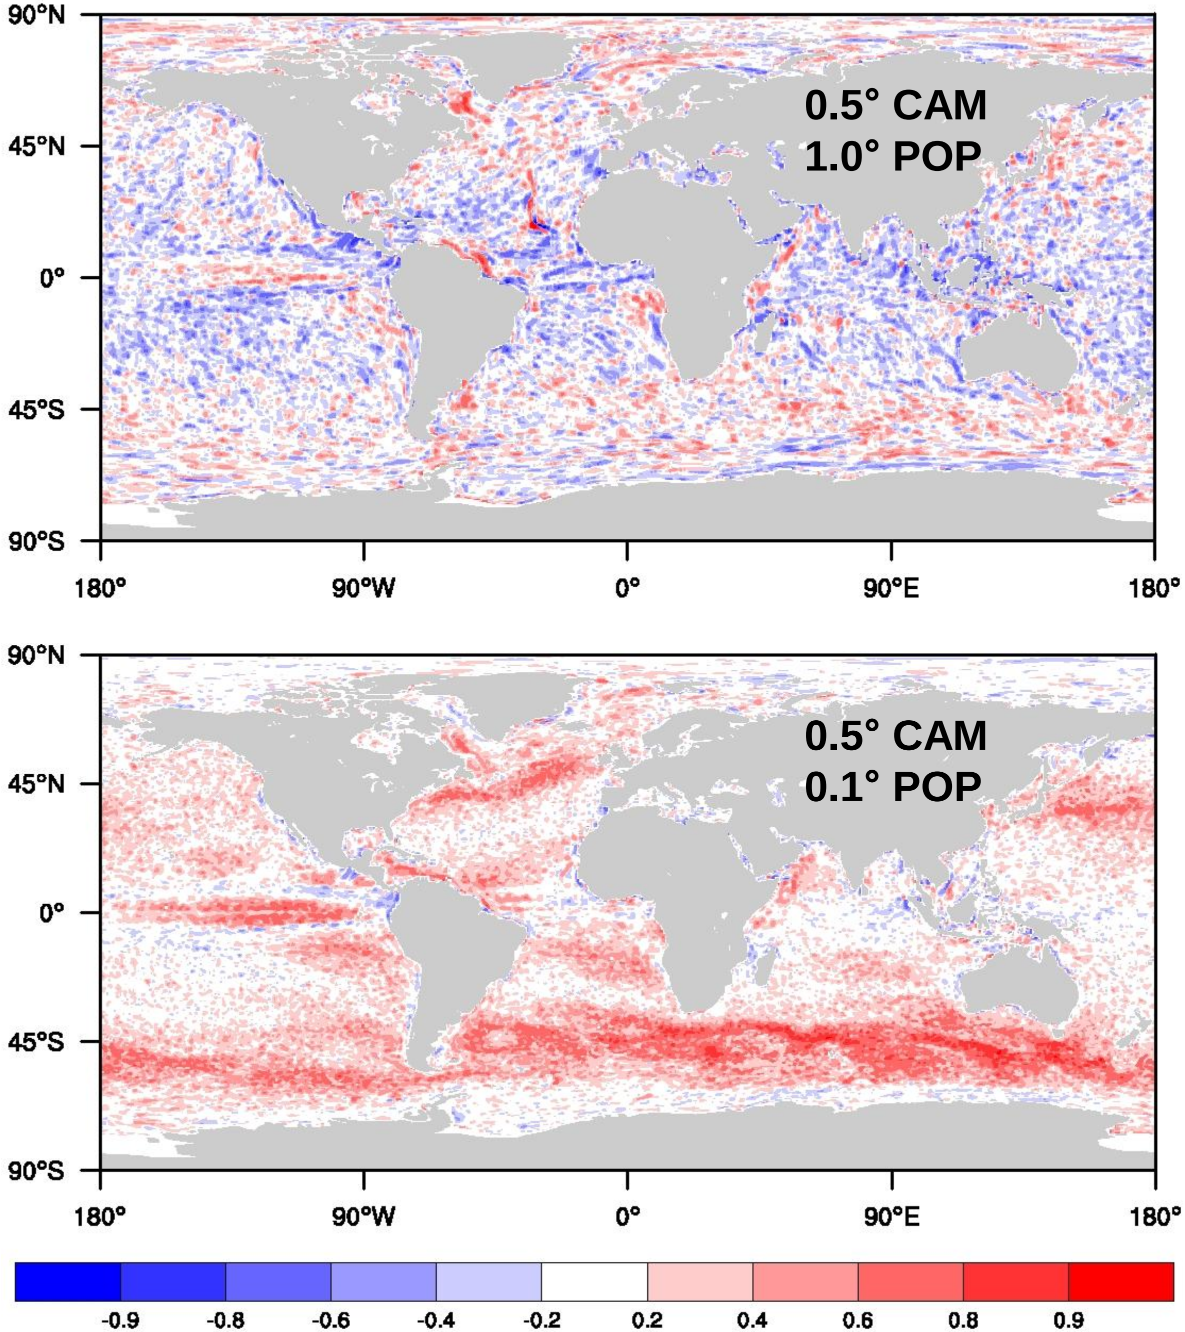

Next-generation global coupled climate models will enter a qualitatively different (and more realistic) dynamical regime.

Courtesy Fran Bryan, NCAR

What About Computing?

Computing Capability

8. It is possible that, within the next 6 years, high-end computers will have

- **peak capability of 100 petaflops (10^{17} floating point operations per second)**
- **Tens of millions of computing units**

Aggregate Computing Capability Still Rising

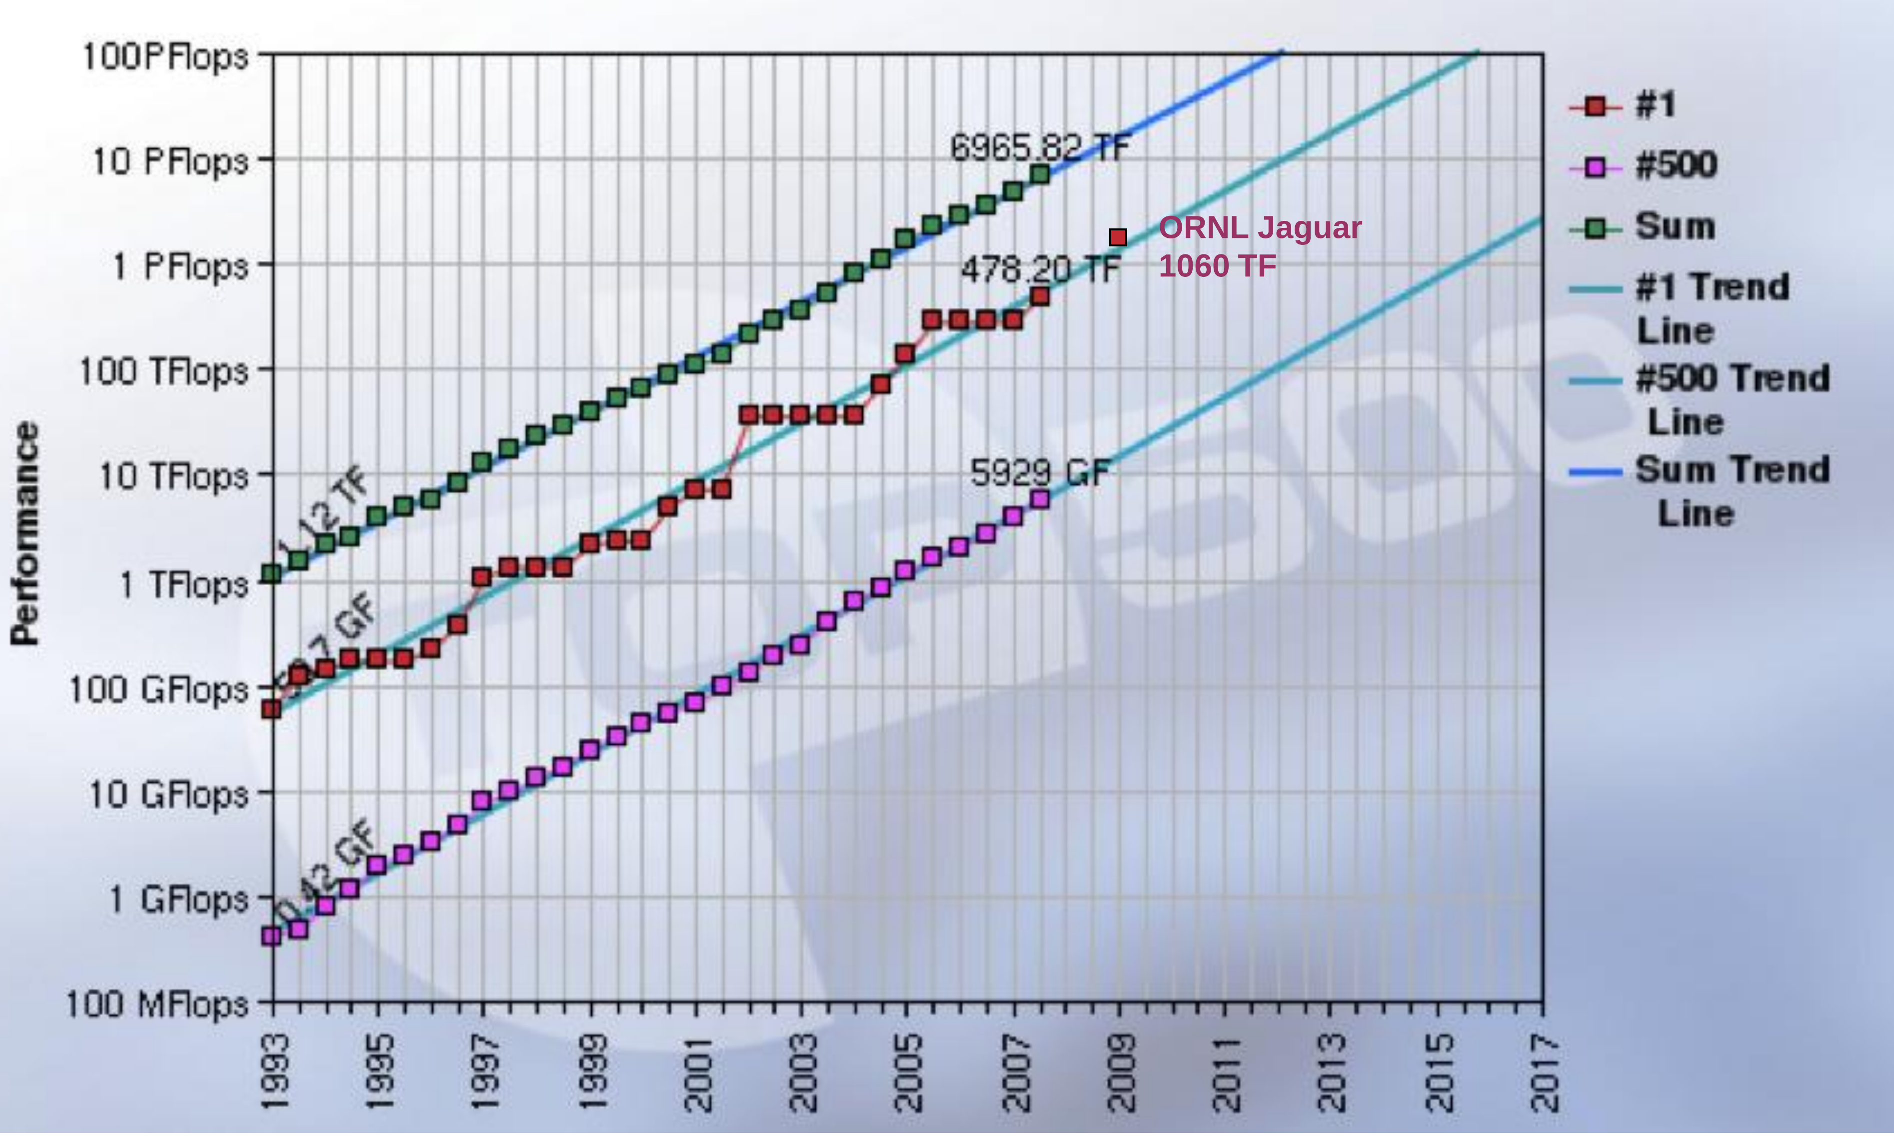

Feature Size Decreasing: -14% yr^{-1}

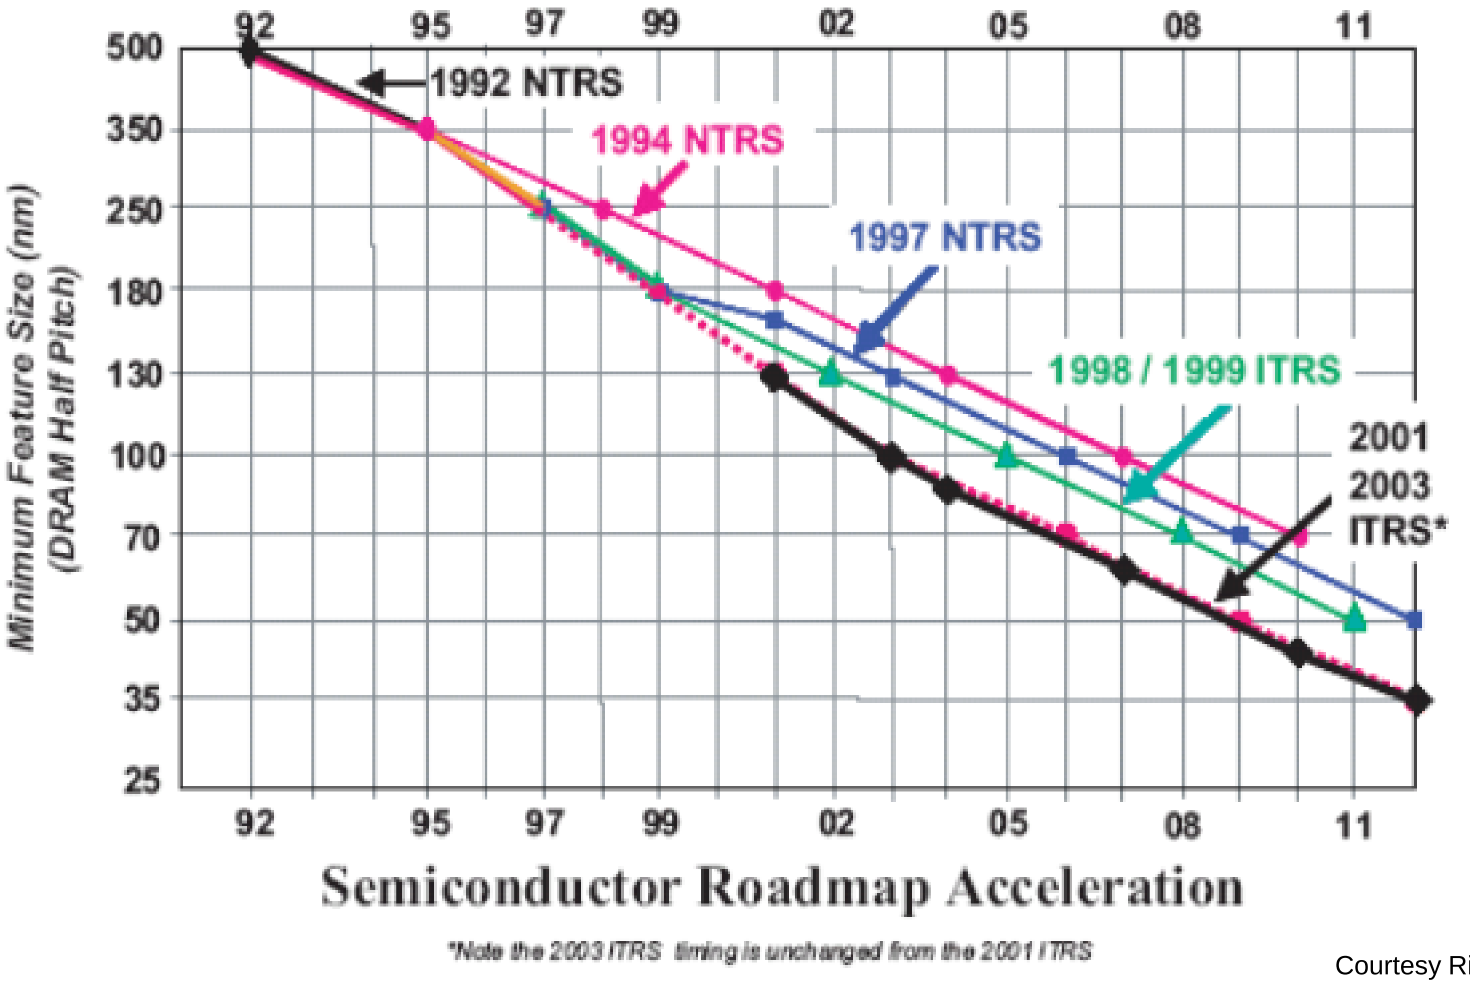

Courtesy Rich Loft, NCAR

Clock Speed - Flattening

- Chip density increasing @ 2X per 2 years
  - Clock speed is not
  - Cores/chip doubling instead
- Hidden parallelism - little or none remaining potential
- $\rightarrow$ Parallelism via software

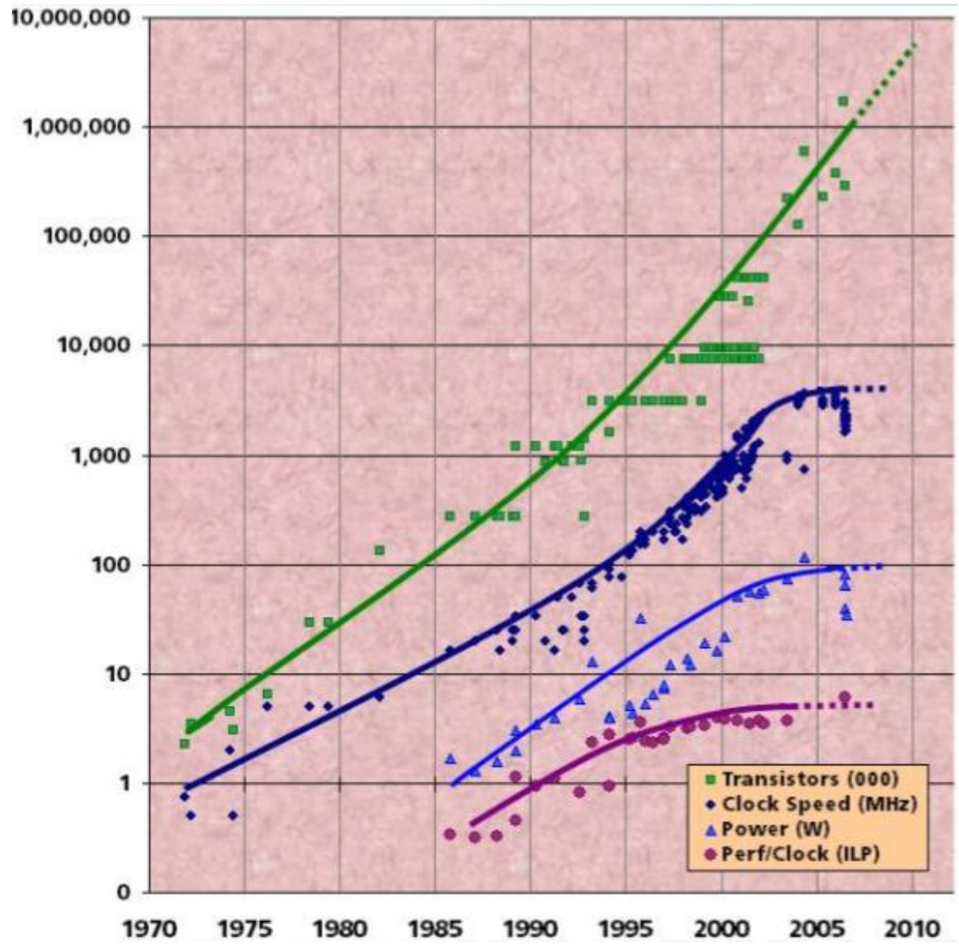

Courtesy Rich Loft, NCAR

Power Density Rising Relentlessly

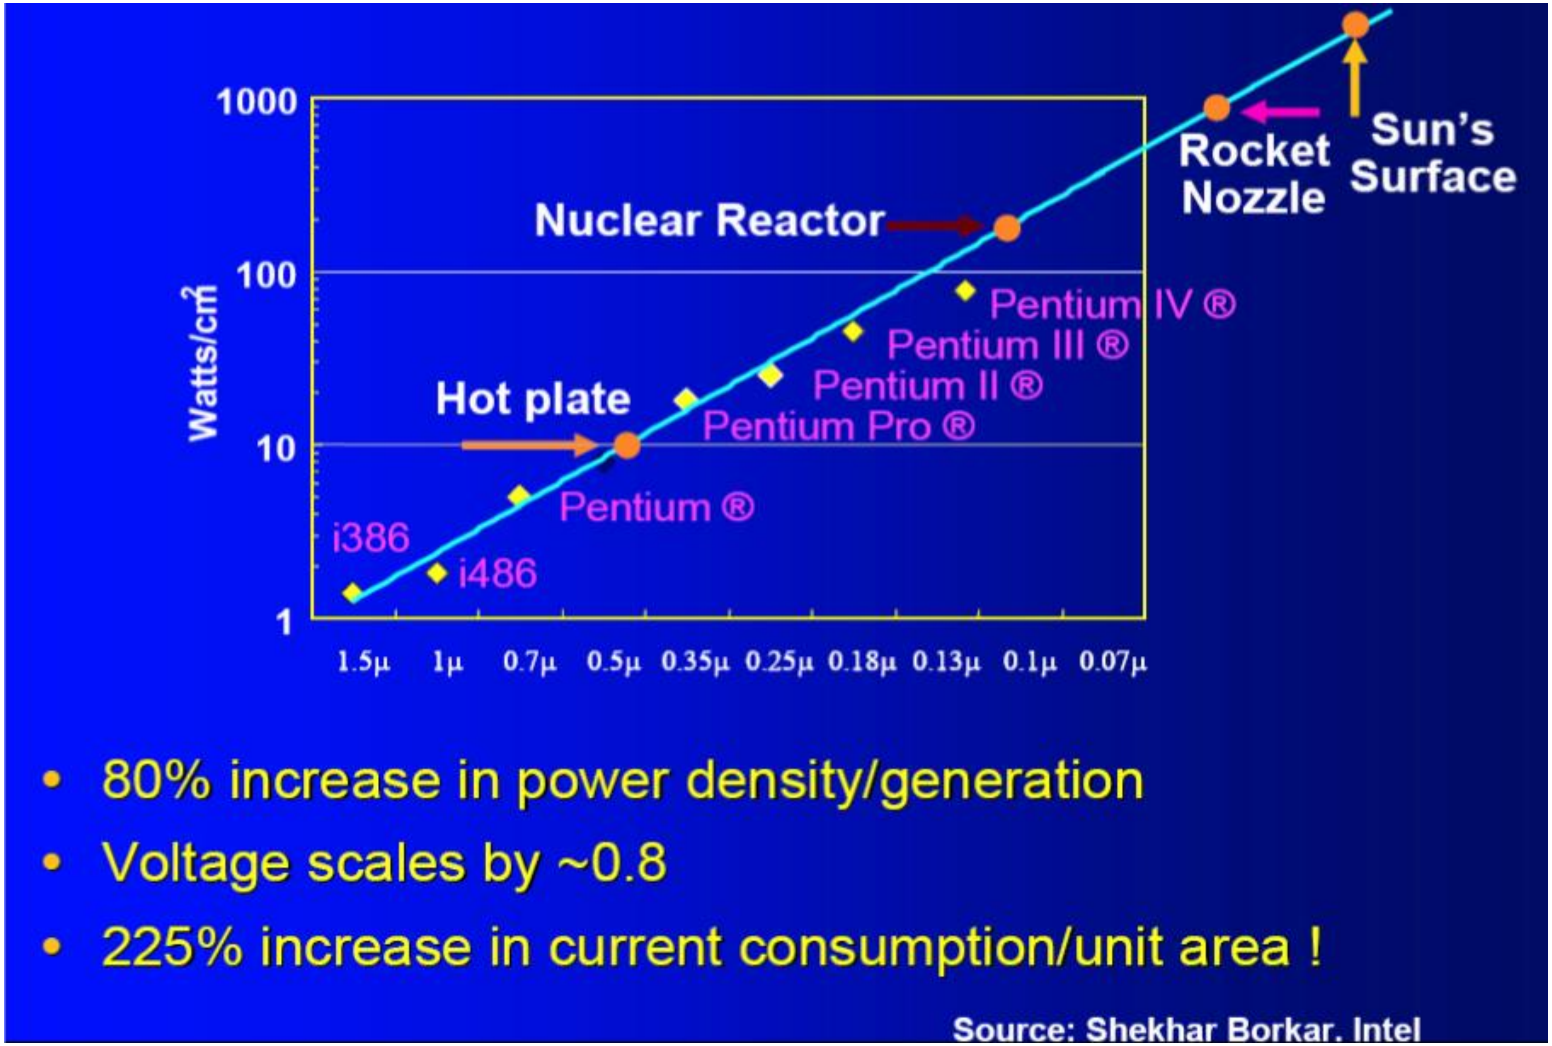

Memory: Hitting a Wall

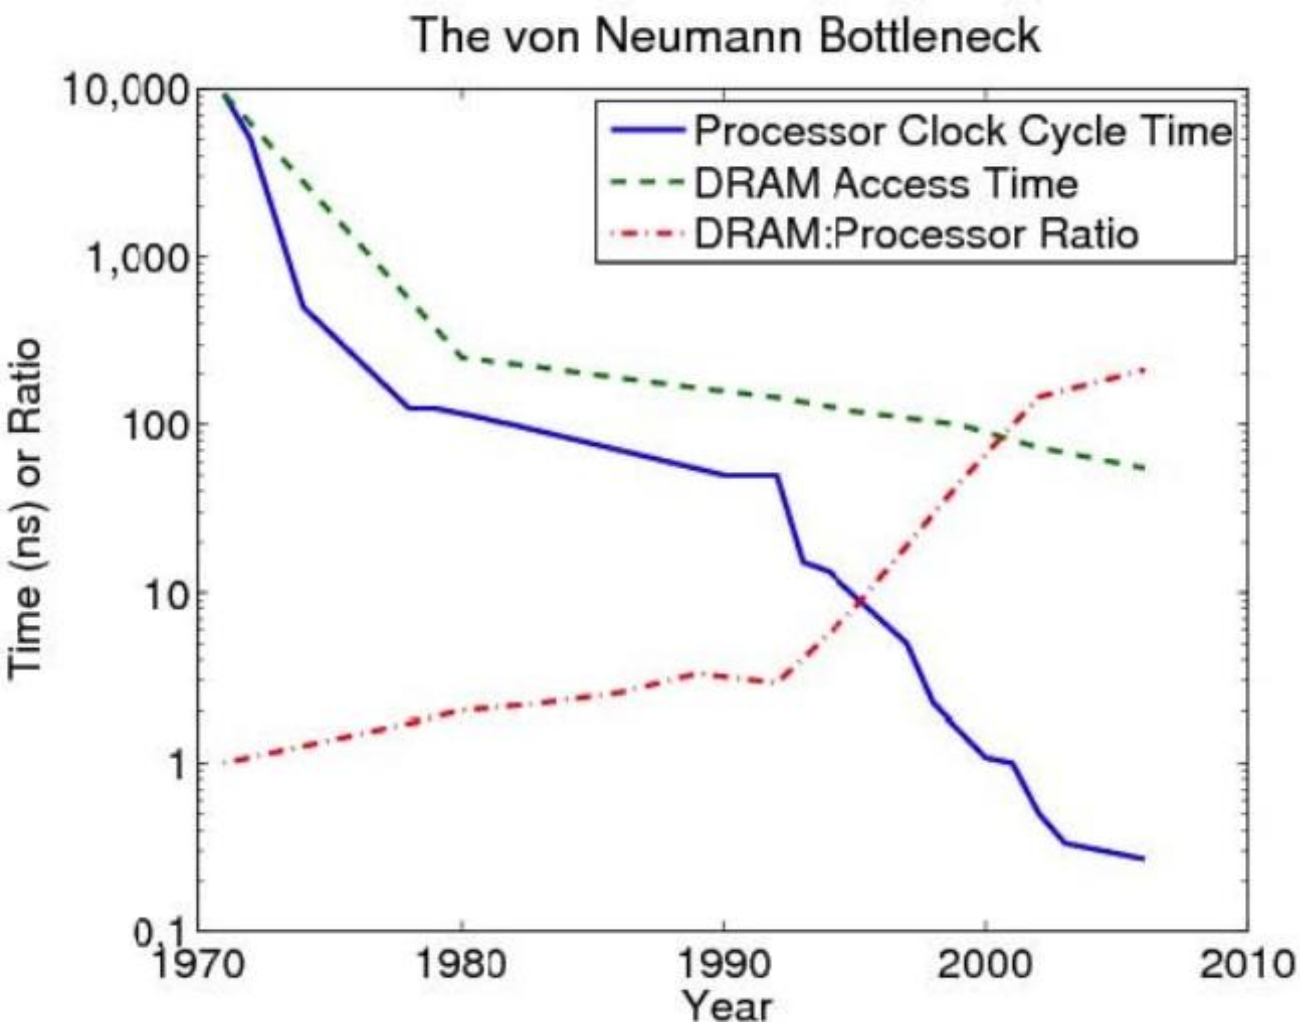

Courtesy Rich Loft, NCAR

Understanding Failures in Petascale Computers

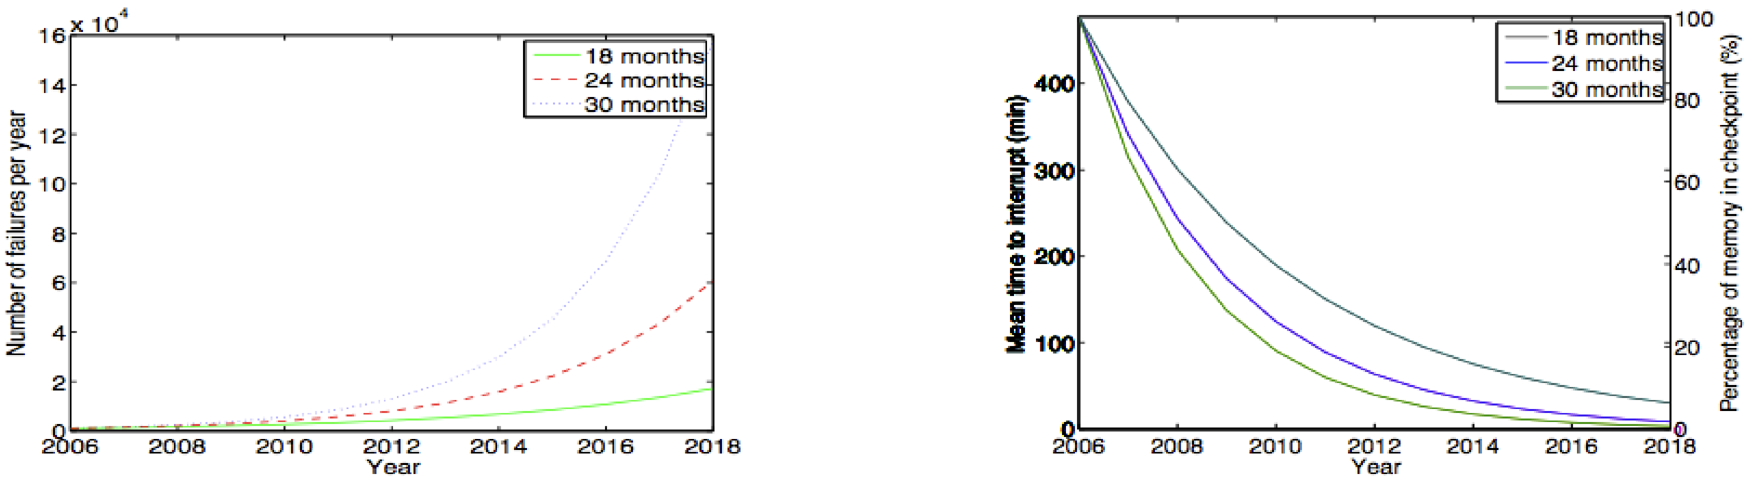

Figure 4. The expected growth in failure rate (left) and decrease in MTTI (right), assuming that the number of cores per socket grows by a factor of 2 every 18, 24 and 30 months, respectively, and the number of sockets increases to stay on top500.org.

Schroeder and Gibson, J. Comp. Phys. (2007)

Understanding Failures in Petascale Computers

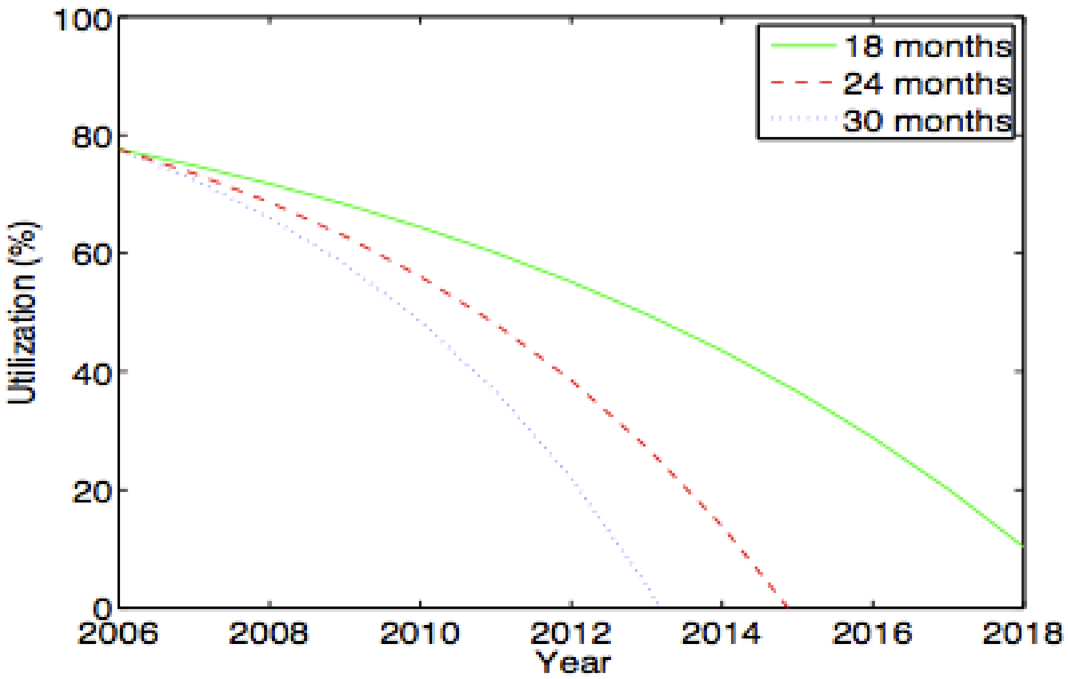

Figure 5. Effective application utilization over time.

Schroeder and Gibson, J. Comp. Phys. (2007)

Understanding Failures in Petascale Computers

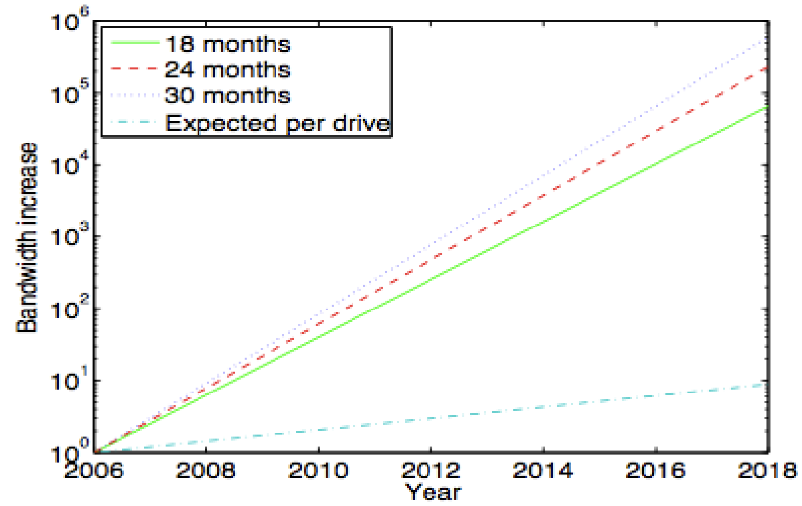

Figure 6. The growth in bandwidth necessary to make checkpoints cheap enough to compensate for the increased failure rate due to growth in number of sockets, assuming the number of cores per socket grows by a factor of 2 every 18, 24 and 30 months, respectively.

Schroeder and Gibson, J. Comp. Phys. (2007)

Current State of (US) Petascale Systems

- Los Alamos Roadrunner (IBM) - 1.6 PFLOPS
  - 16K Opteron cores
  - 128K Cell SPEs
- LLNL BlueGene (IBM) - 1.0 PFLOPS
  - 65K powerPC cores
- ORNL Jaguar (Cray XT5) - 1.0 PFLOPS
  - 150K Opteron cores (Barcelona quads) *
- NASA Ames Pleiades (SGI Ice) - 0.6 PFLOPS
  - 51K Xeon cores (Harpertown and Clovertown quads)
- NICS Kraken (Cray XT5) - 0.6 PFLOPS
  - 66K Opteron cores (Barcelona quads) *
- NCSA Blue Waters (IBM) - 10 PFLOPS (2011)
  - 320K Power7 cores (octacores?)

* Being upgraded to Istanbul hexa-cores in 2009

Kraken XT-5 at NICS

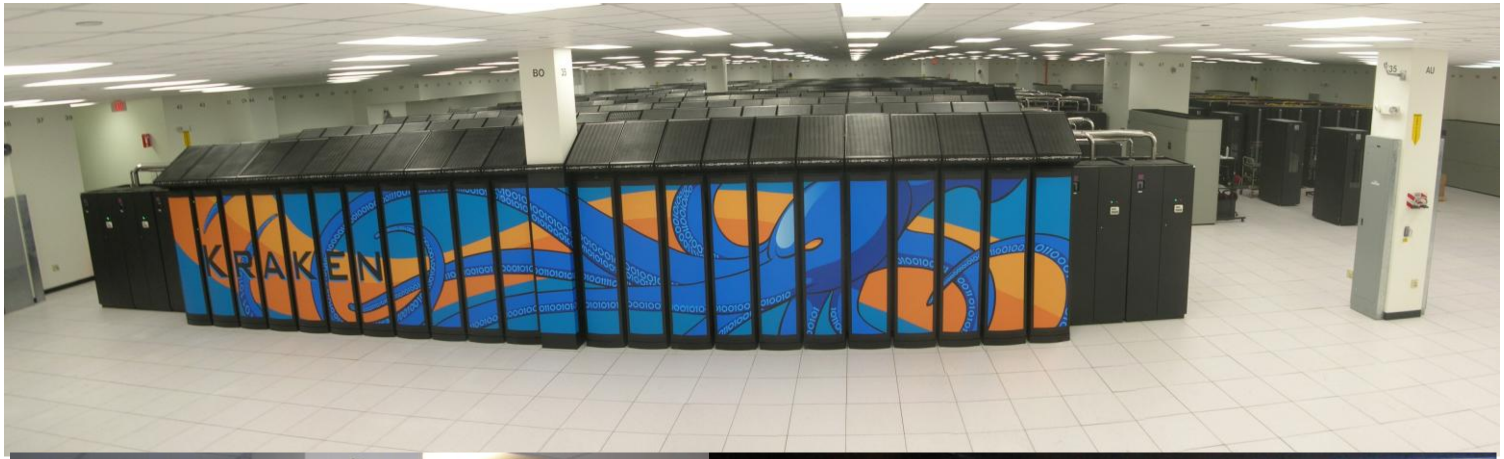

Courtesy of Pat Kovatch, NICS

Likely Platform in 2015

- **22 micron feature size (Moore's "law")**
- **Heterogeneous architecture**
  - **~32 million heavyweight processors**
  - **~8 million light-weight processors (e.g. GPU)**
  - **Two decades experience with communicating-sequential-processes execution model and message-passing programming model --> <u>community is ill-prepared for new era of heterogeneous multi-core platforms</u>**
- **16 PB memory**
- **100 PFLOPs peak performance**
  - **75% sustained on high-performance Linpack**
  - **What percentage will fluid dynamics codes sustain?**
- **22 MW power requirement**

Courtesy Thomas Sterling, LSU

Balancing Future Demands on Computing Power

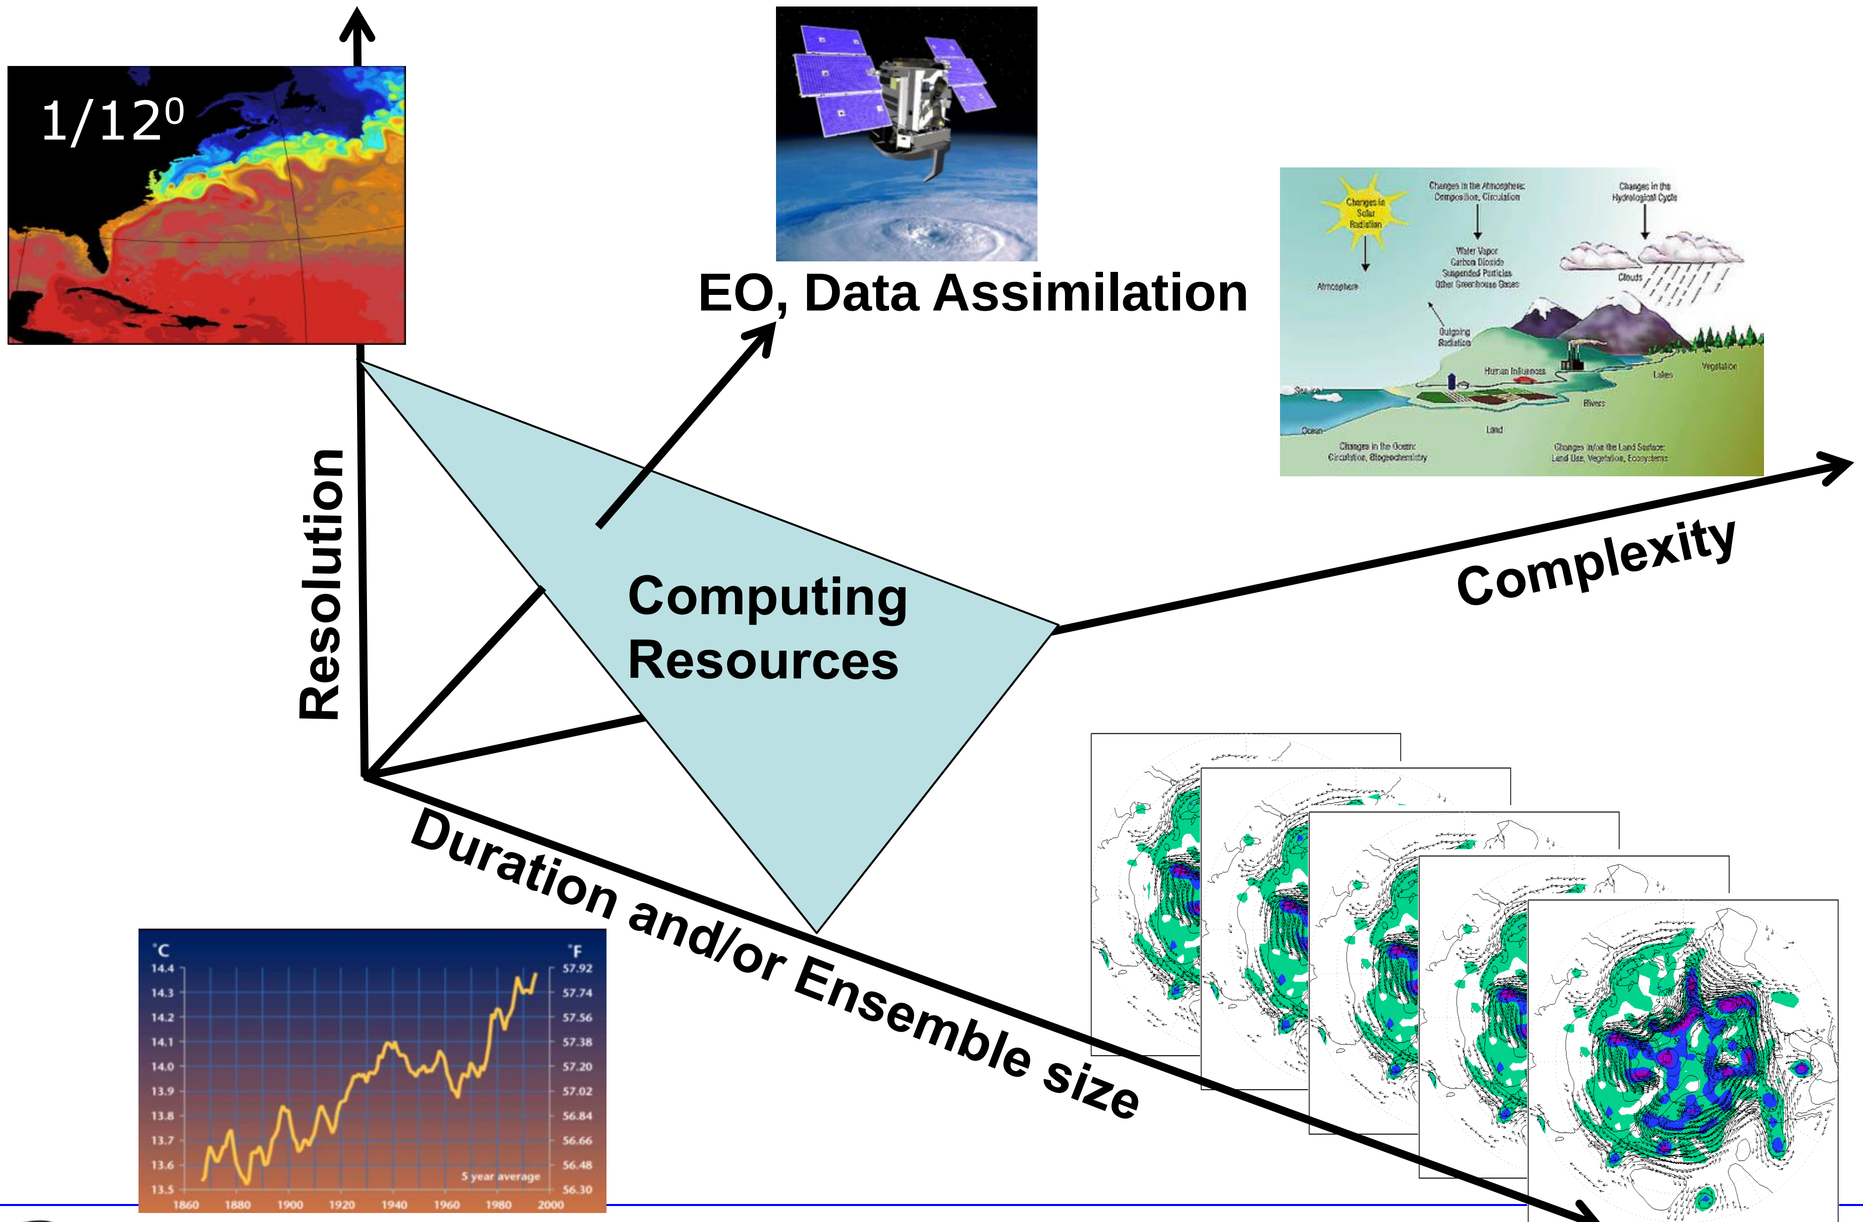

New Computing Paradigm

9. The unprecedented level of parallel computation and heterogeneous computing platforms will require new systems software, new algorithms and substantial re-coding of existing climate models.

10. The massive scale of these computations will also require new methods in workflow management, data management, and visualization.

Static, Weighted Load Balancing Example:
Space Filling Curves for CICE4 @ 1° on 20 processors

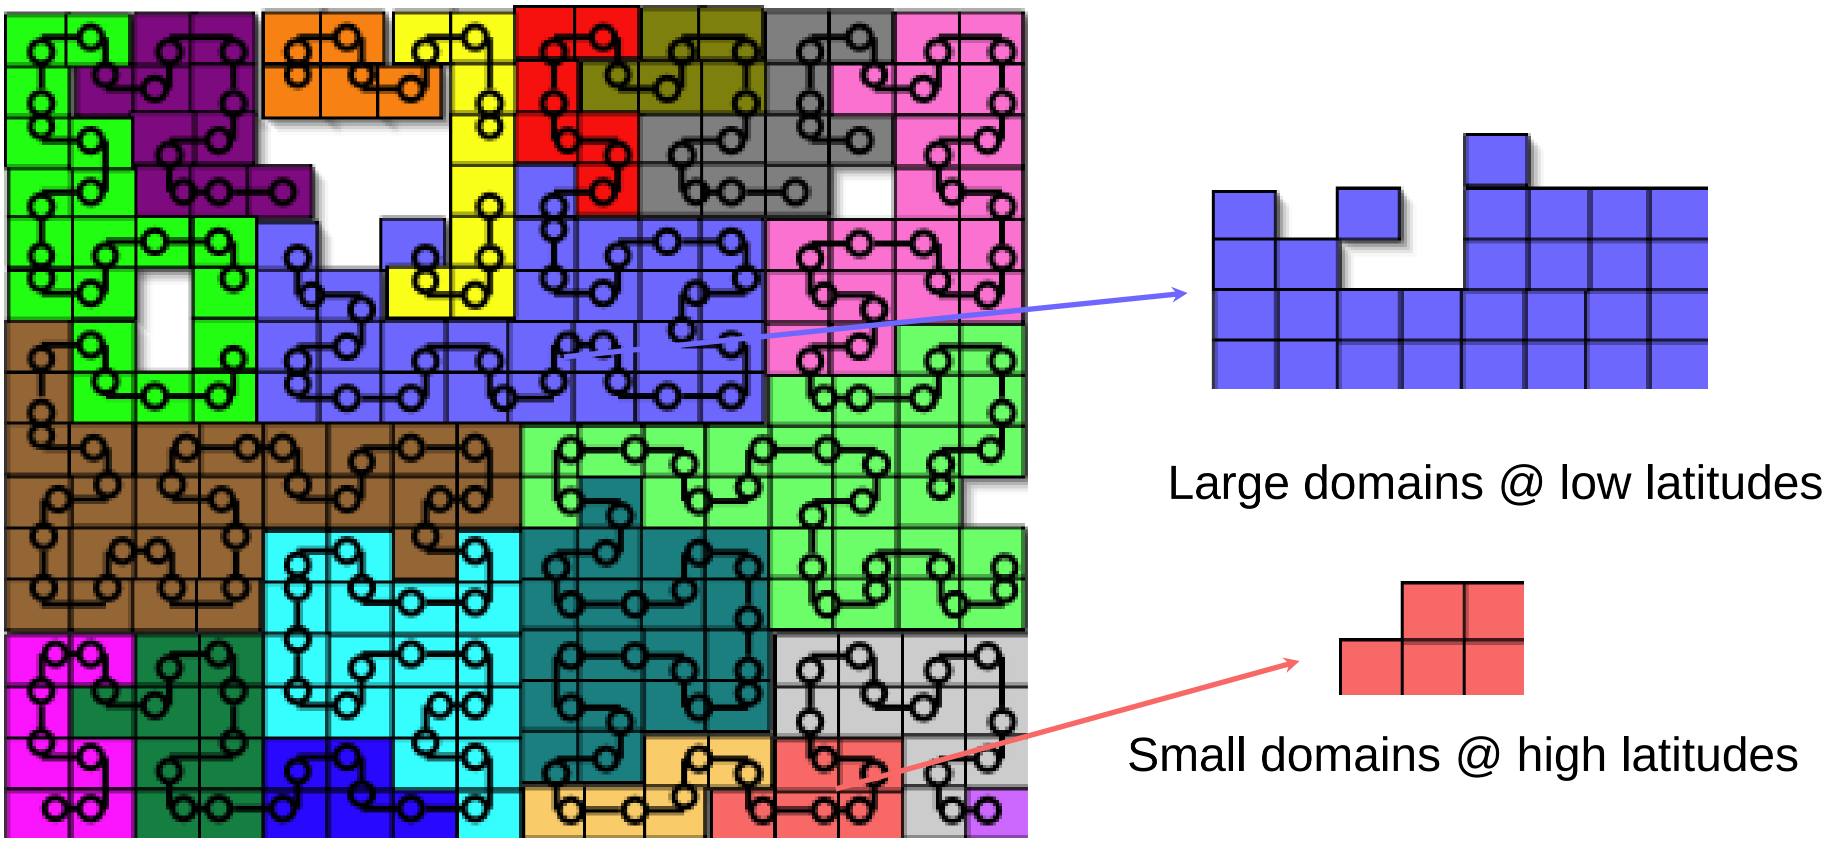

Courtesy John Dennis, NCAR

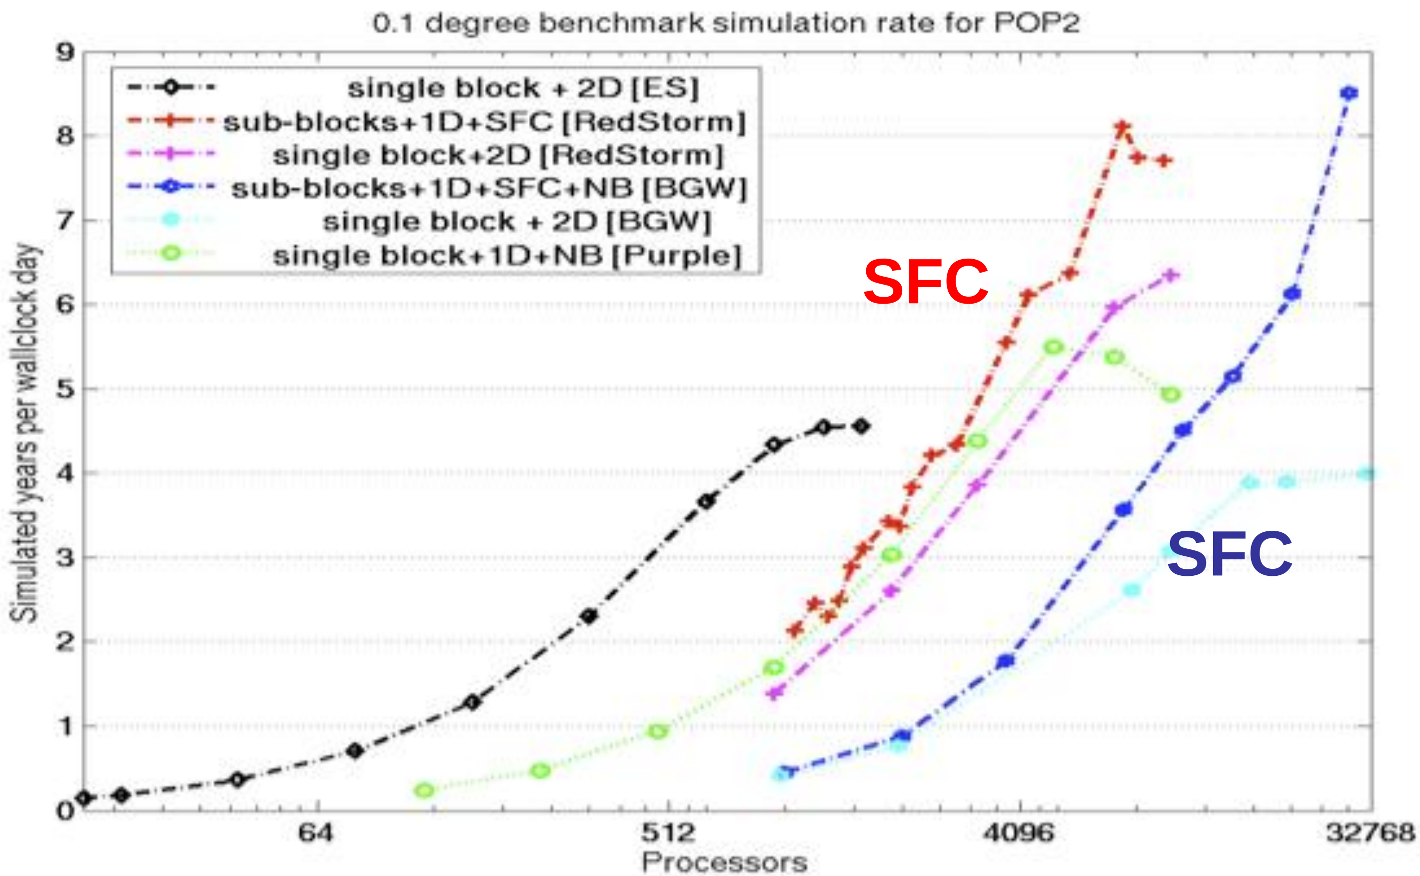

Credit: J. Dennis, Y. Yoshida, M. Taylor, P. Worley

22 October 2007

Example of highly scalable algorithms: High Order Multiscale Modeling Environment (HOMME)

- Sphere is decomposed into 6 identical regions using a central projection (Sadourny, 1972) with equiangular grid (Rancic et al., 1996).
- Avoids pole problems, quasi-uniform.
- Non-orthogonal curvilinear coordinate system with identical metric terms

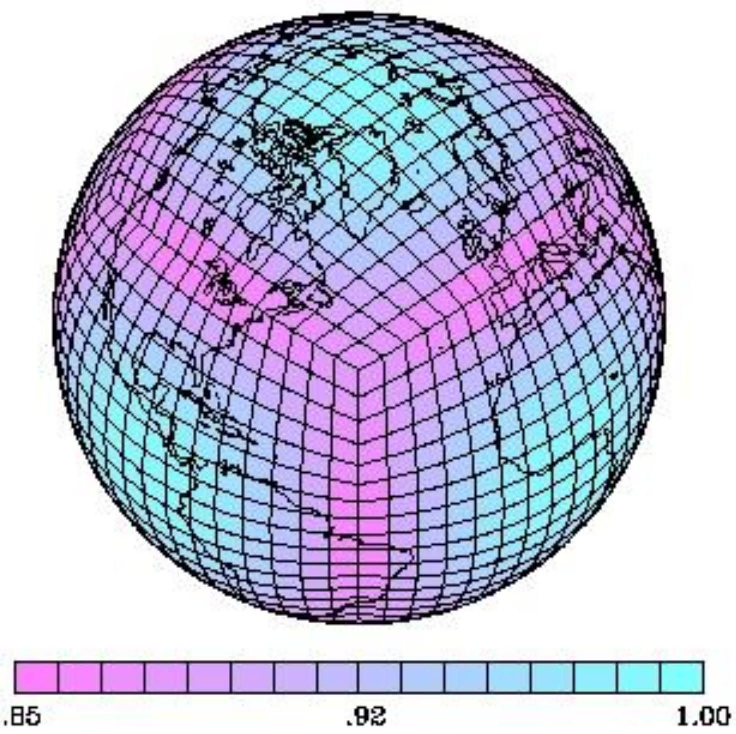

Ne=16 Cube Sphere Showing degree of non-uniformity

Today, scalable dycores can't break 1 simulated year per day even at 12.5 km…

Full CAM Physics/HOMME Dycore on Blue Gene/L
Parallel I/O library used for physics aerosol input and input data
(work COULD NOT have been done without Parallel IO)
Work underway to couple to other CCSM components

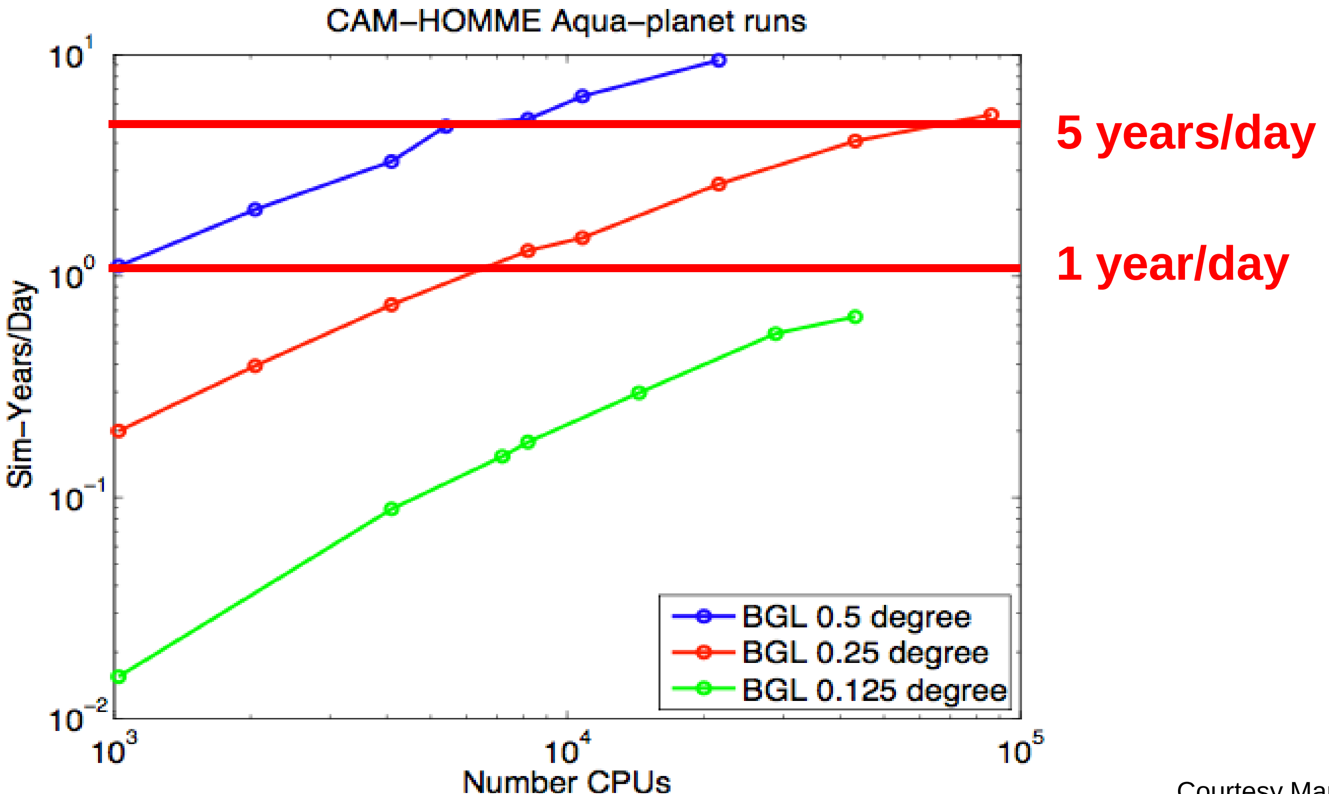

Courtesy Mark Taylor, SNL

The Dynamical Core Ain't Everything …

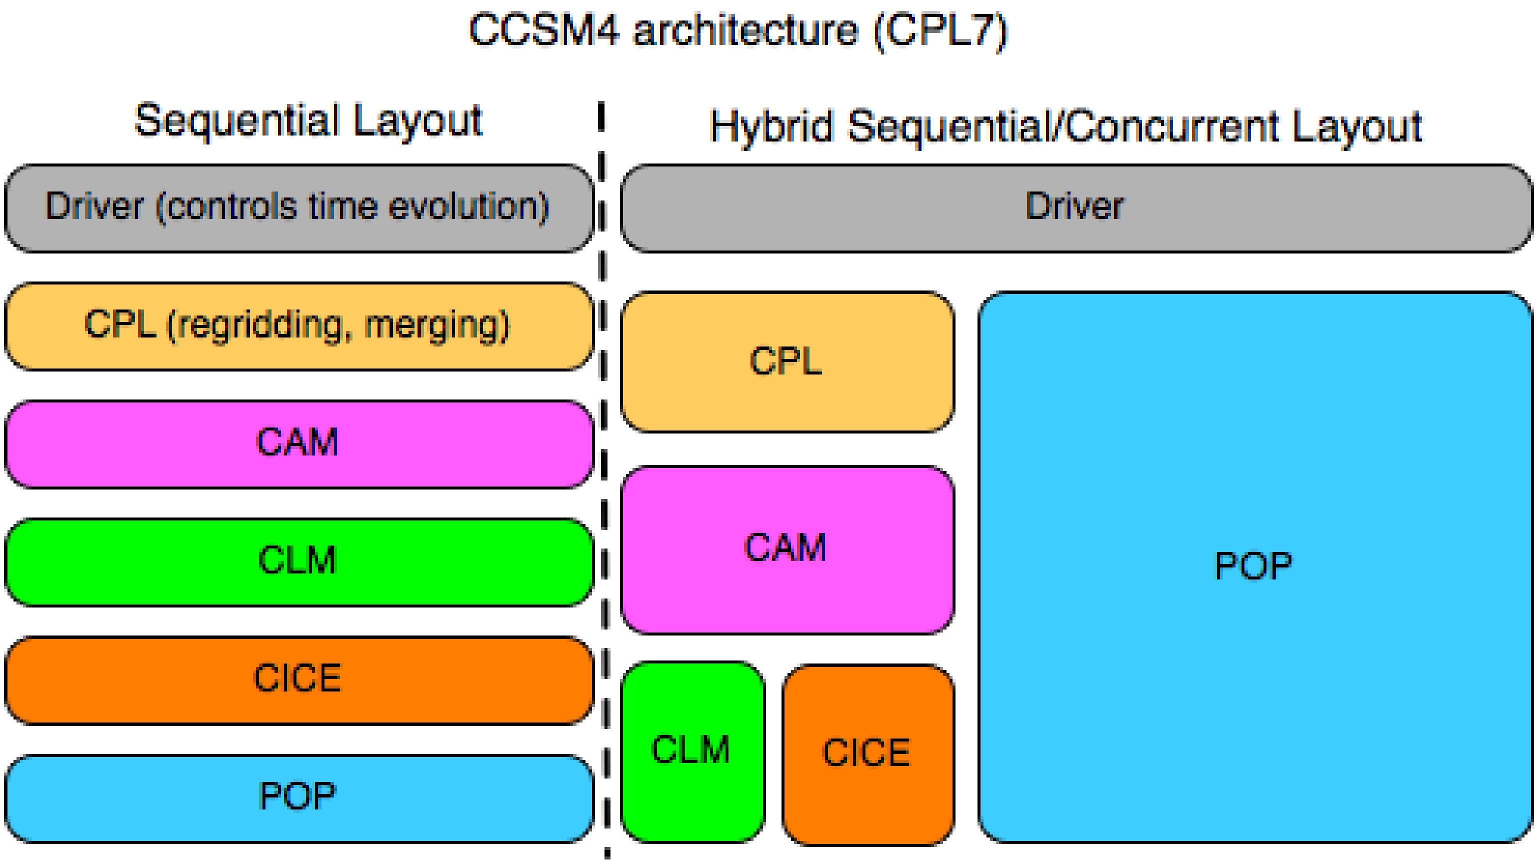

Courtesy Mariana Vertenstein, NCAR

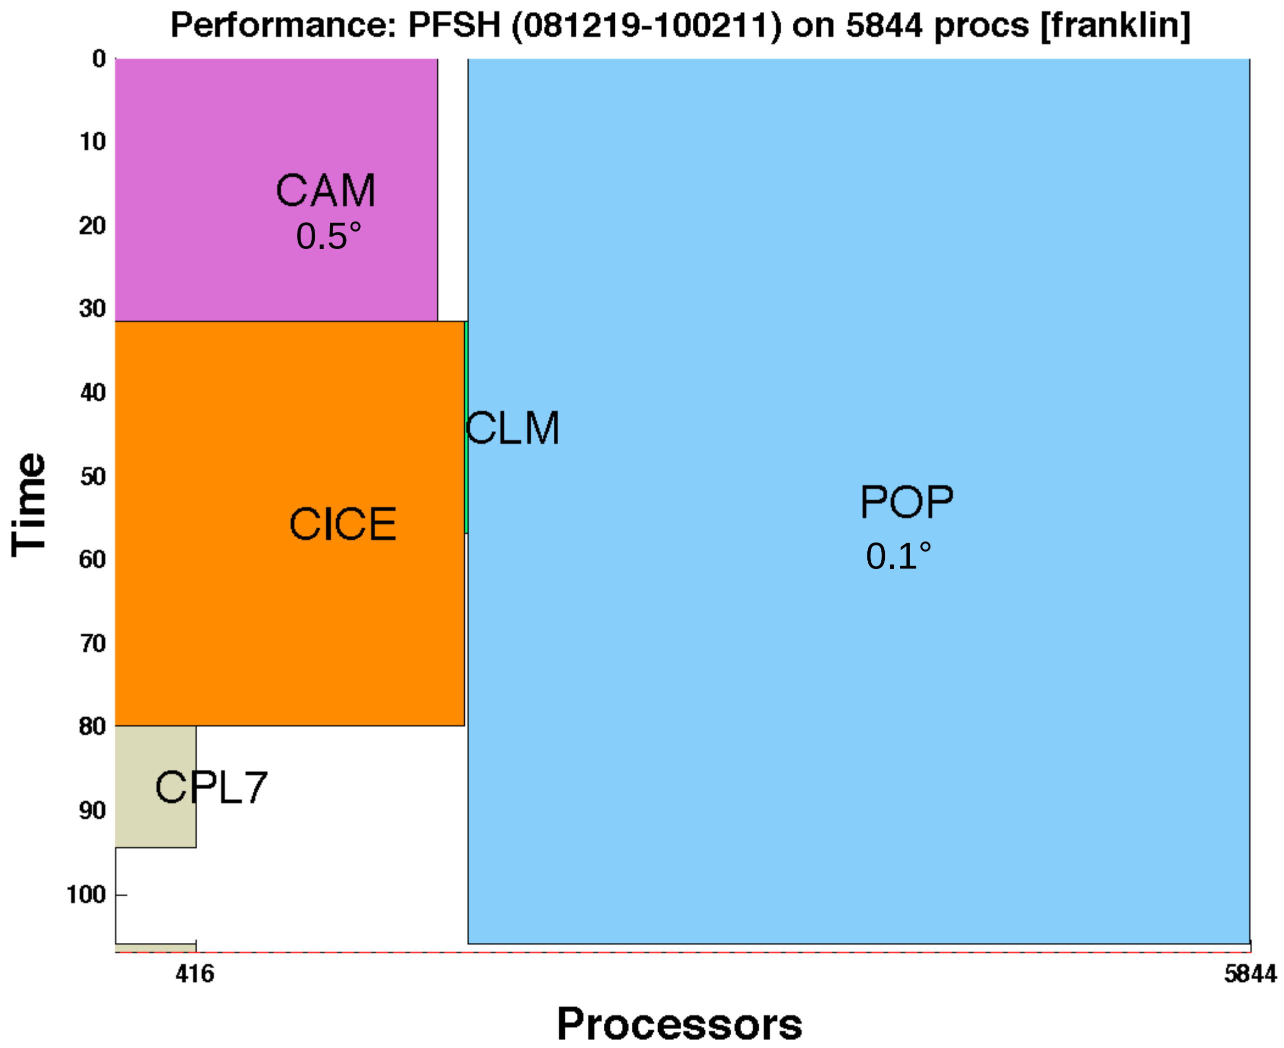

Courtesy John Dennis, NCAR

PetaApps

NSF PetaApps proposal funded - three-year research and development effort, aimed at enabling a broad climate science capability for petascale systems

- Interactive ensembles using CCSM/CPL7
- Incorporate and examine use of PGAS language (Titanium) in CCSM

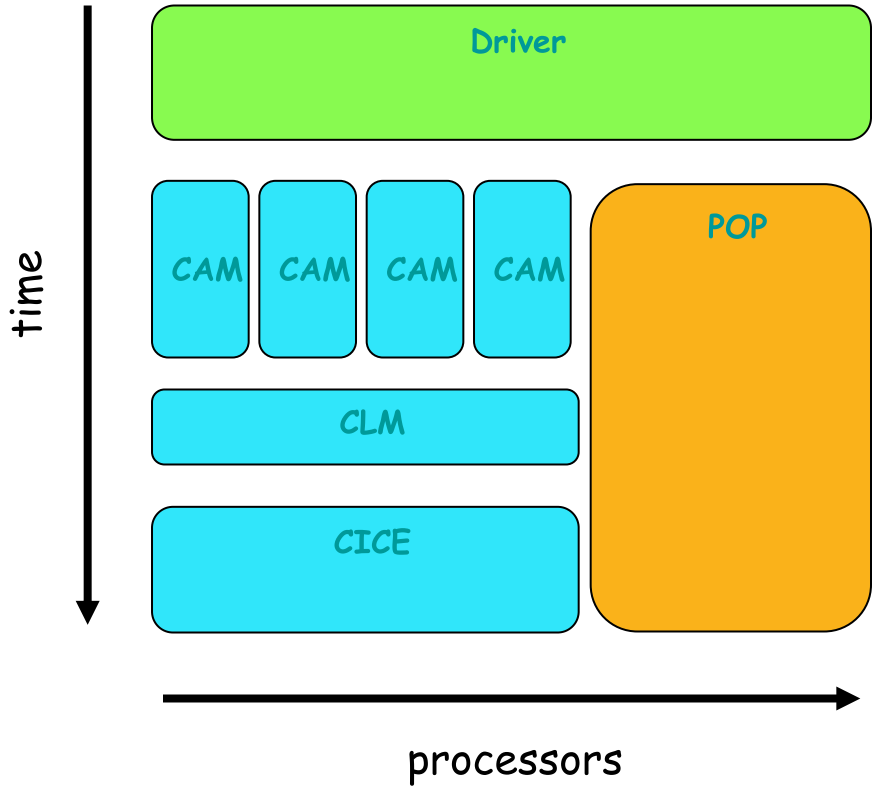

Interactive ensembles will be used to understand how oceanic and atmospheric weather noise impacts climate variability

Titanium will be used in the flux coupler to examine impact of PGAS language approach to performance and memory footprint

Courtesy Mariana Vertenstein, NCAR

Global 1-km Model - A Simple How-To Guide

- **Objective**: Global, atmospheric dycore (100 levels) simulating 1 yr per wallclock-day (satisfying CFL)
- **Flops**: Need 6.37 PFLOPS sustained for 1 km resolution
  - Fluid codes on microprocessor-based computers get 5% of peak (e.g. 1 GFLOPS per core on IBM Power6)
  - Need ~6 million processors (assuming quite optimistic scaling) in 180X180X180 torus topology
  - Power requirement: 200-300 MW
- **Communications**: Need link bandwidths ~100 GB/sec = 10X Cray XT5 link bandwidth
  - ~510M horizontal grid points; patch-size 9x9x100 per core for 6 million cores
  - For 10 prognostic variables, 35 MB/sec/patch keeps up with 1 Gflops per core, assuming overlapping communications. Without overlapping communications, need 10-20X = 350-700 MB/sec/patch.
- **Data Storage**: 100 sim-years = 1 exabyte generated in 100 days
  - Need high levels of data management parallelism
  - Visualization requires sophisticated data compression and subsetting techniques.
  - Other multipliers: ensembles, non-atmospheric components, …

Courtesy Rich Loft, NCAR

Upshot …

- To do 1 km modeling at 1 year per day in 2009 would require at least a 6 million core machine and 200 MW.
- That is beyond current engineering capability
  - Processor speed is stuck
  - Amdahl's Law is a formidable barrier
  - Mean time between failures goes like 1/N for N components
  - How long does it take to synchronize 10 million threads?
- Dynamical timestep goes like N^{-1}
  - CFL is merciless
  - The cost of (subcycled) dynamics relative to physics increases with resolution (N)
    - e.g. if dynamics takes 20% at 25 km it takes 86% of the time at 1 km
- The amount of data produced is O(exabytes) and vastly exceeds the capabilities of serial postprocessing systems to analyze

Courtesy Rich Loft, NCAR

The Good News …

- To do **20 km** modeling at **1 year per day** would require a dedicated machine with **~15,000 cores**
- That is within reach of current engineering capability

- To do a **global change run at 20 km**, for which **5 years per day** would be a more reasonable throughput requirement, it is probably necessary to have a dedicated machine with **~100,000 cores**
- Scaling will be a challenge on current architectures, but Blue Waters may rise to this challenge

Algorithmic Acceleration is Important: The History of MHD Simulations

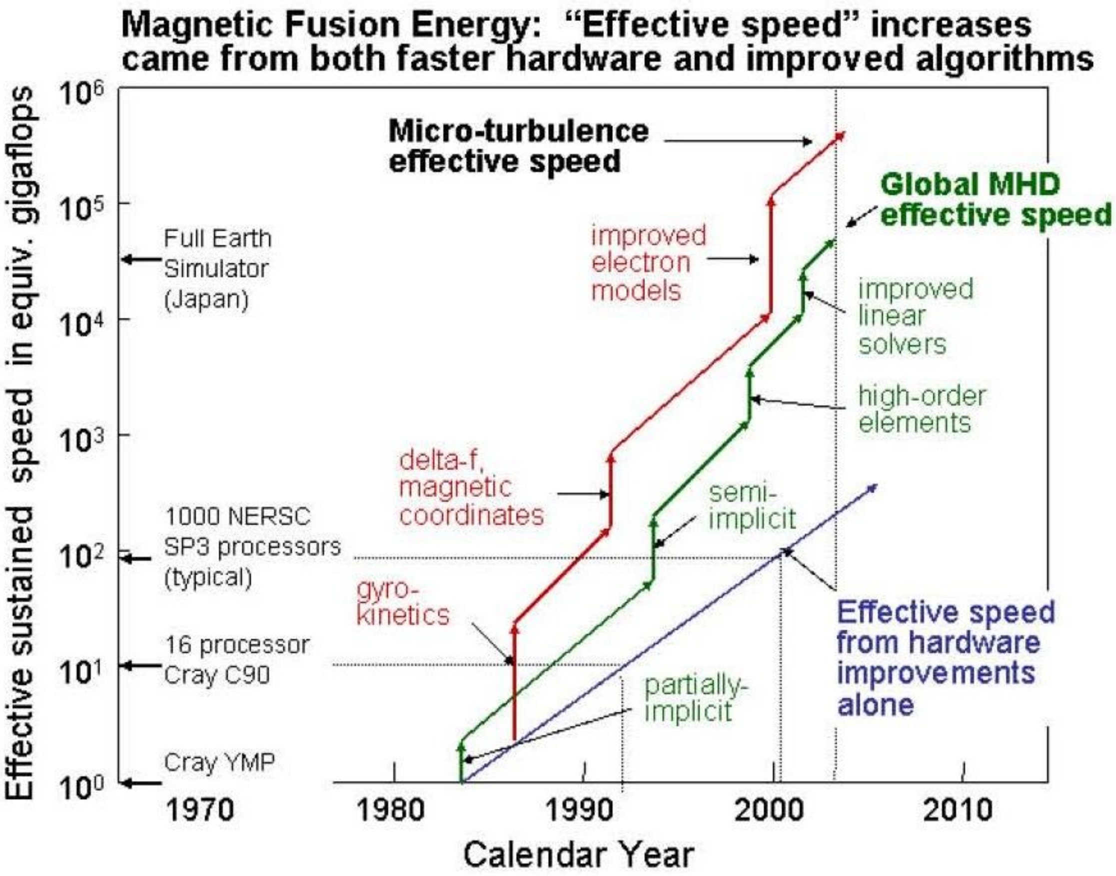

SCaLeS Report, 2004

Trade-offs in Numerical Properties

Desired Properties

- Local and global conservation
- High order accuracy
- Computational efficiency
  - Flops per unit space-time volume
- Geometric properties
  - Flexibility
  - Uniformity
- Non-oscillatory advection - monotonicity
- High scalability

Issues

- Conservation schemes force specific PE formulations
- High order accuracy conflicts with computational efficiency and monotonicity (Gudonov)
- Non local/irregular communication and memory access patterns of SLT, AMR, etc.
- Solver non-locality (e.g. of inner products) conflicts with scalability.

Courtesy Rich Loft, NCAR

Existing computers have high penalties for non-locality

HPC Challenge Benchmarks

HPL

System	Achieved (TFLOPS)	Affiliation
Cray XT5	902.0	ORNL
IBM BG/L	259.2	LLNL
IBM BG/P	191.3	ANL

GFFT

System	Achieved (TFLOPS)	Affiliation
IBM BG/P	5.080	ANL
Cray XT3	2.870	SNL
Cray XT5	2.773	ORNL

Best HPL to Best FFT Ratio: 177.6 !

Computational Intensity (CI)

- Compute Intensity:

  CI = Total Operations/(Input + Output data)
- GFLOPS = CI*Bandwidth
- Bandwidth expensive, flops cheap
- The higher the CI, the better we're able to exploit this state of affairs

Courtesy Rich Loft, NCAR

NASA High-End Computing (HEC) Program Study
NASA Goddard, July 24, 2008

Presented at 2008 International Supercomputing Conference (ISC '08), June 18-20, Dresden, Germany.
Study lead: Shujia Zhou, Reference: http://www.hec.nasa.gov/news/features/2008/cell.074208.html

- Workload: Solar Radiation component from Goddard Earth Observing System Model, Version 5 (GEOS-5) – 20% of GEOS-5 computing time.
  - Computationally intensive & I/O time smaller than computation tin
  - Relatively small code size (~2,000 Fortran lines)
  - Vertical column data allowed for independent computation, reducing communication
- Porting
  - Ported code to C. ***Note: IBM XL Fortran compiler for Cell is nc available***
  - Adding Direct Memory Access library calls to communicate with SPE's.
  - Determine best mapping across 8 SPE's per Cell/B.E. processor
- Currently working towards exploiting the Cell Fortran compiler with a hybrid version of the solar radiation component.
- Future work…
  - GEOS-5 atmospheric model component for **turbulence**, **moisture**, **chemistry**, and other processes, which collectively take about 50 percent of GEOS-5 computing time.
  - More sophisticated models, including **cloud-resolving** models – requiring more than 10x computing power

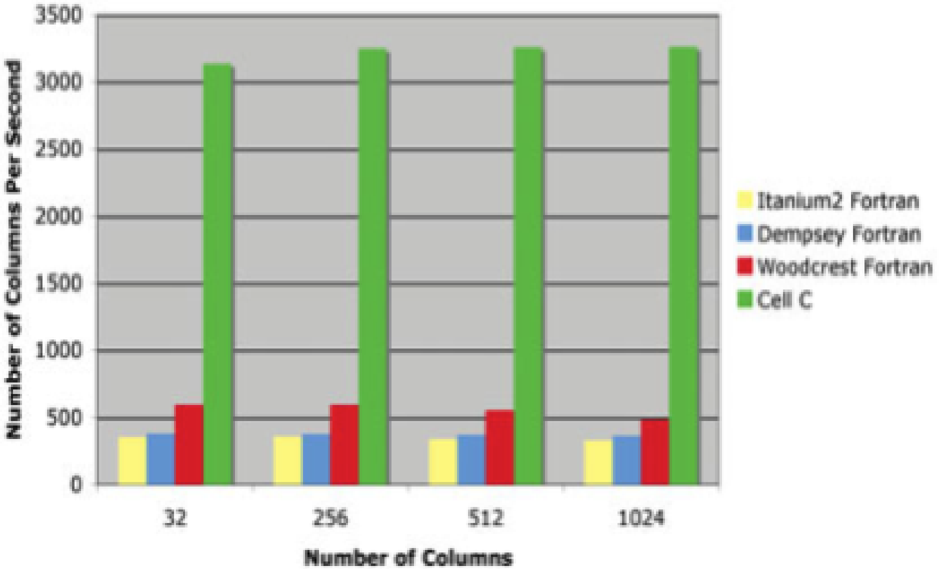

Cell was …
- **6.76x faster than Intel Xeon Woodcrest (2.66 GHz)**
- **8.91x faster than Intel Xeon Dempsey (3.2 GHz)**
- **9.85x faster than Intel Itanium2 (1.5 GHz)**

QS22 is 3.73% faster than QS20 data values above.

Courtesy Kent Winchell, IBM

- Two standards evolving from different sides of the market

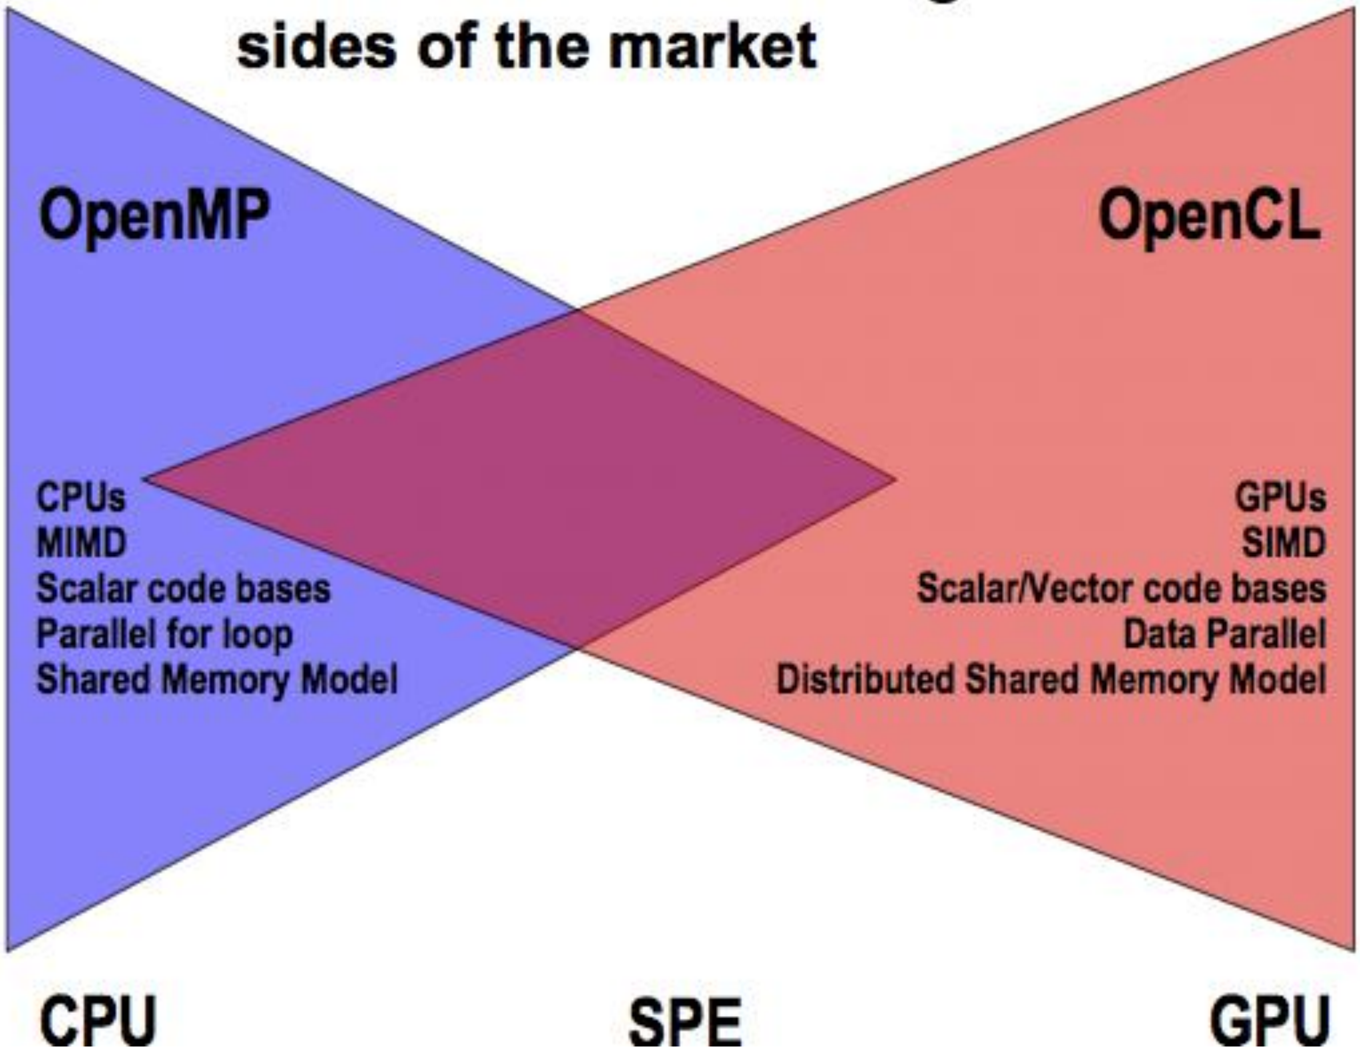

Courtesy Kent Winchell, IBM

Different approaches to exploit hybrid systems

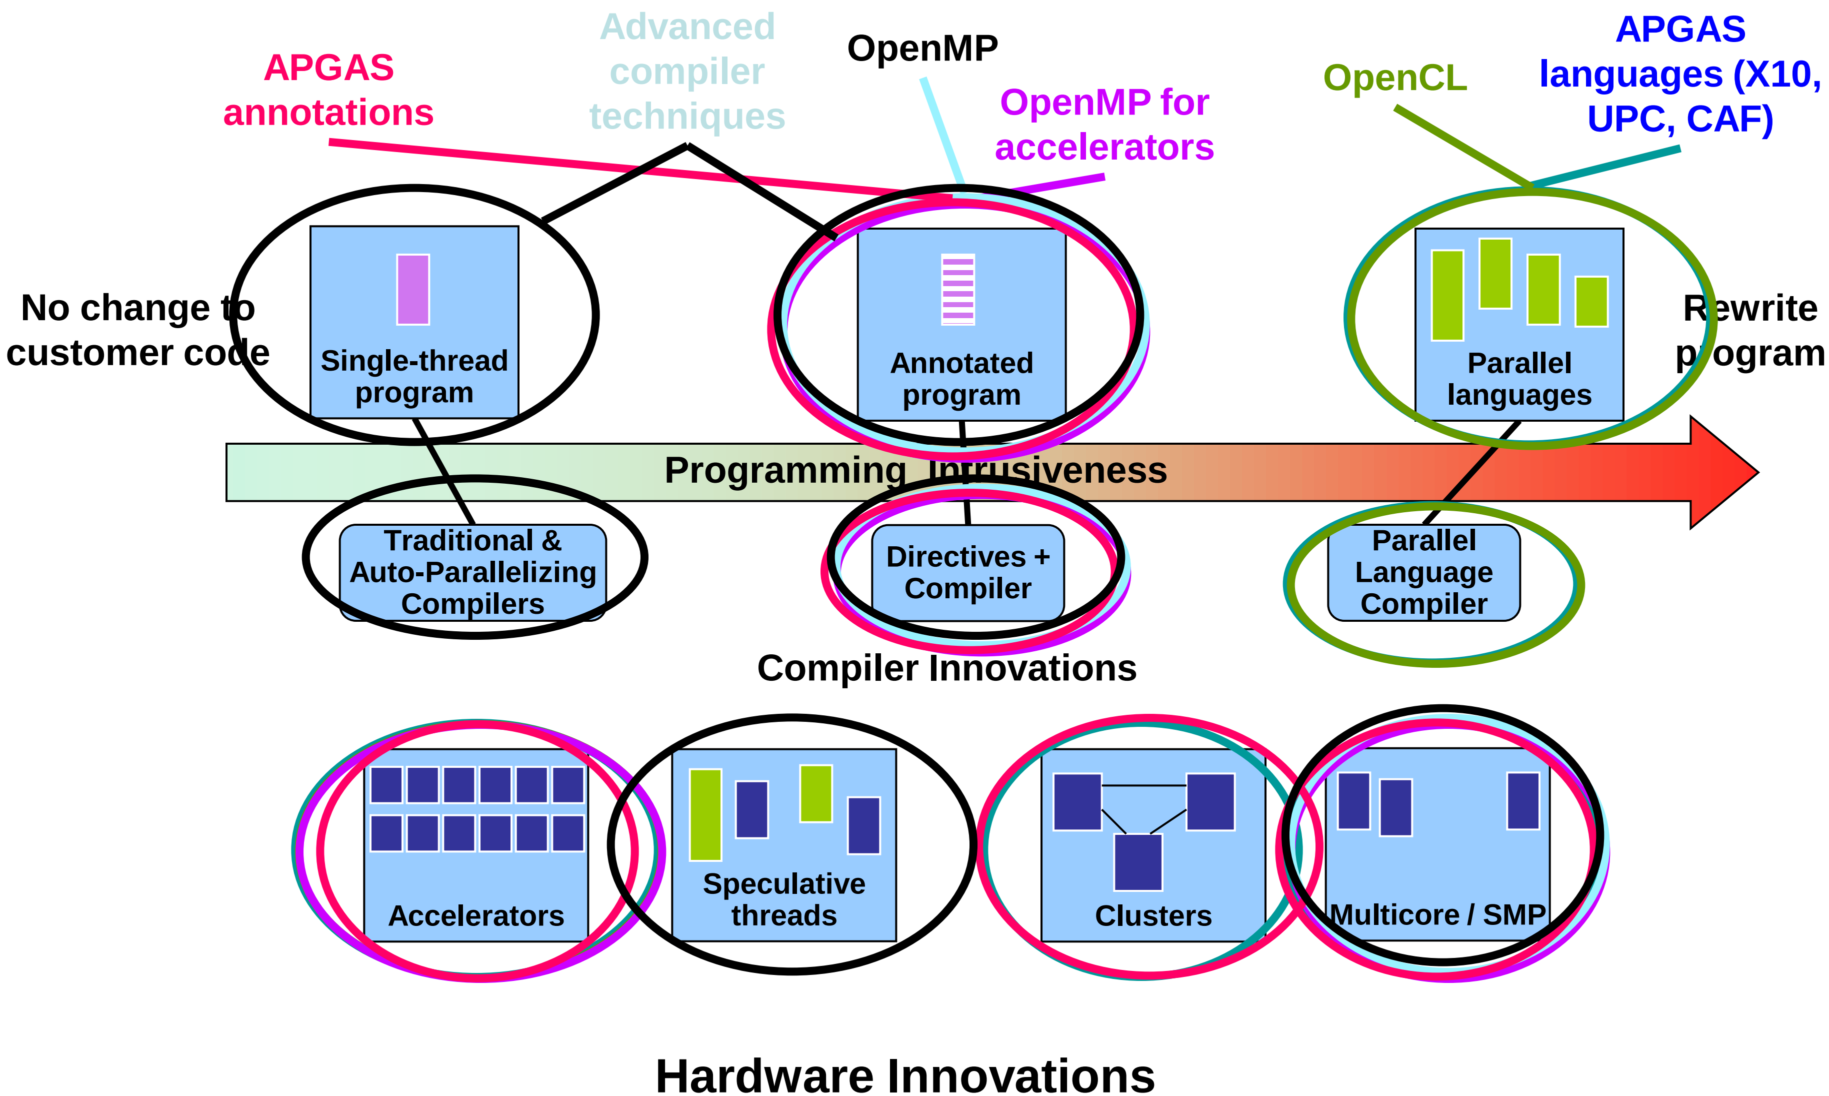

Where's the tipping point for Accelerators: Cell and GPU's?

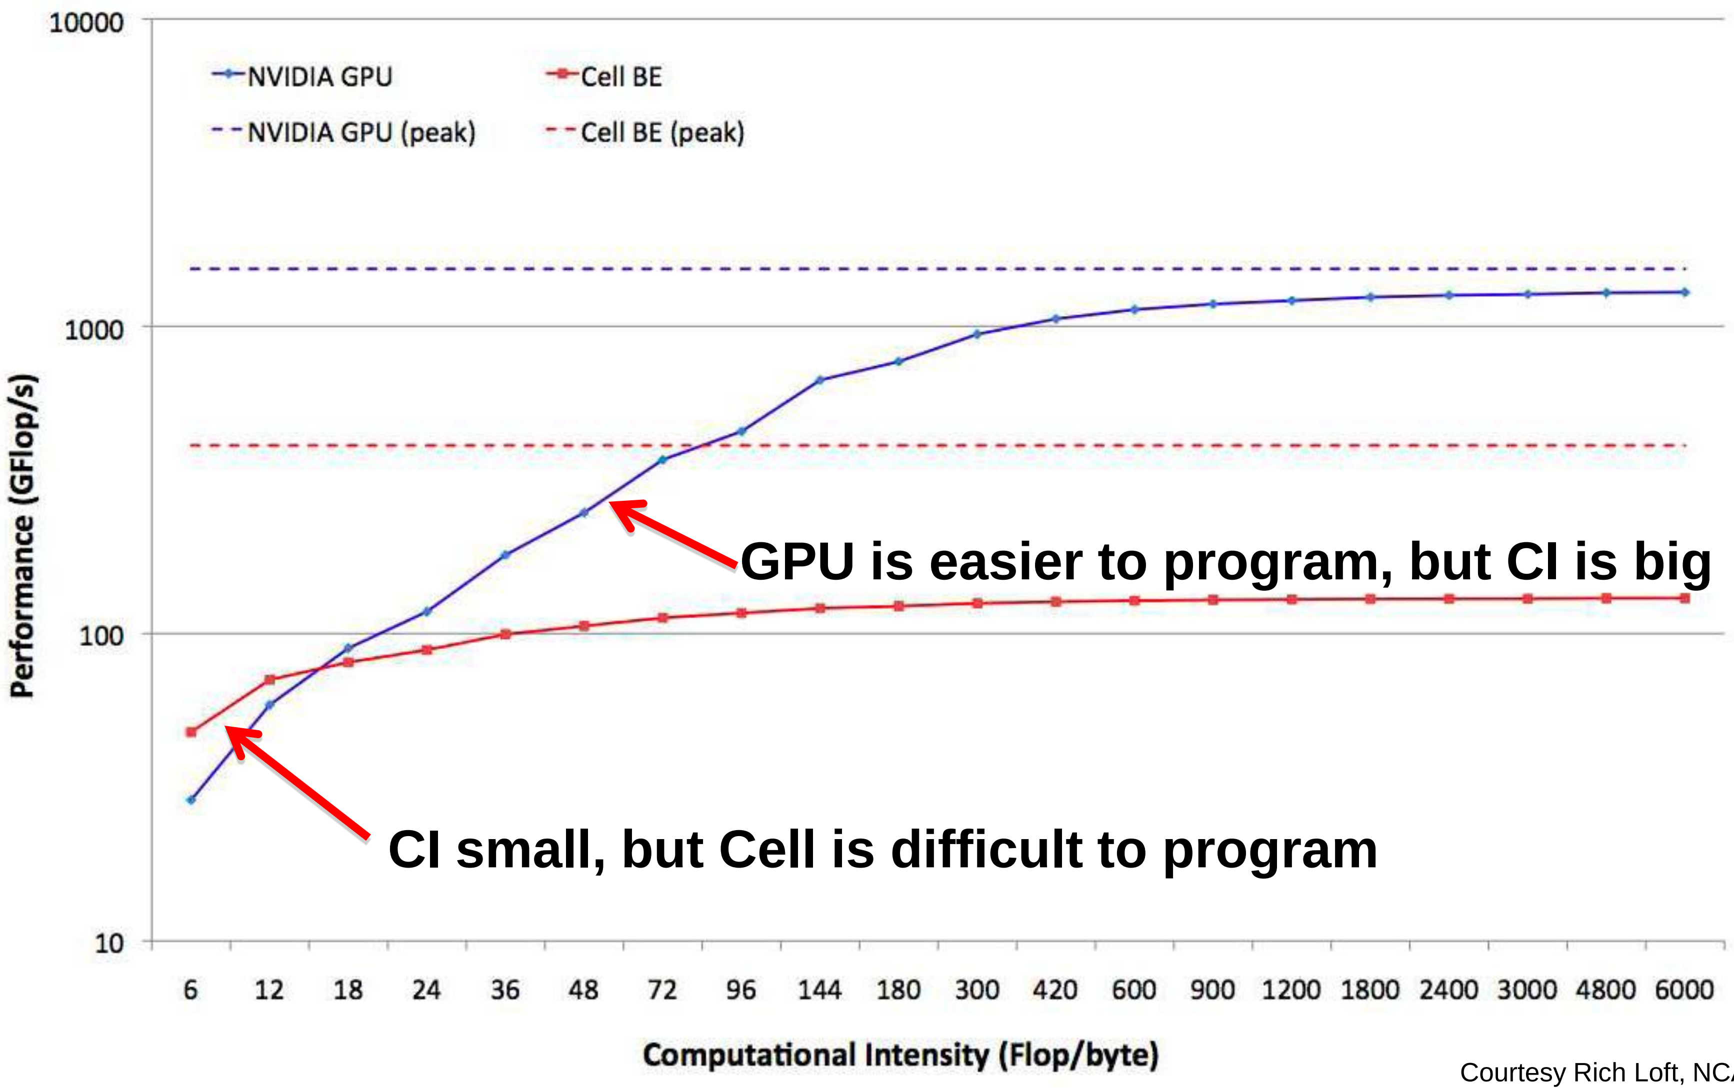

Courtesy Rich Loft, NCAR

Example - Super-parameterization: An intermediate step in the quest for a cloud-system-resolving AGCM

- Super-parameterization is embarrassingly parallel but extremely expensive (~150x traditional physics)

- The computational cost of the following simulations are approximately the same:
  - A millennium-long simulation using a traditional climate model.
  - A few years-long simulation using a traditional climate model with CRCP (MMF or SP)
  - A day-long simulation of a cloud-system-resolving AGCM O(1 km)

Courtesy Rich Loft, NCAR

Running “Super Parameterization” on an Accelerator Cluster

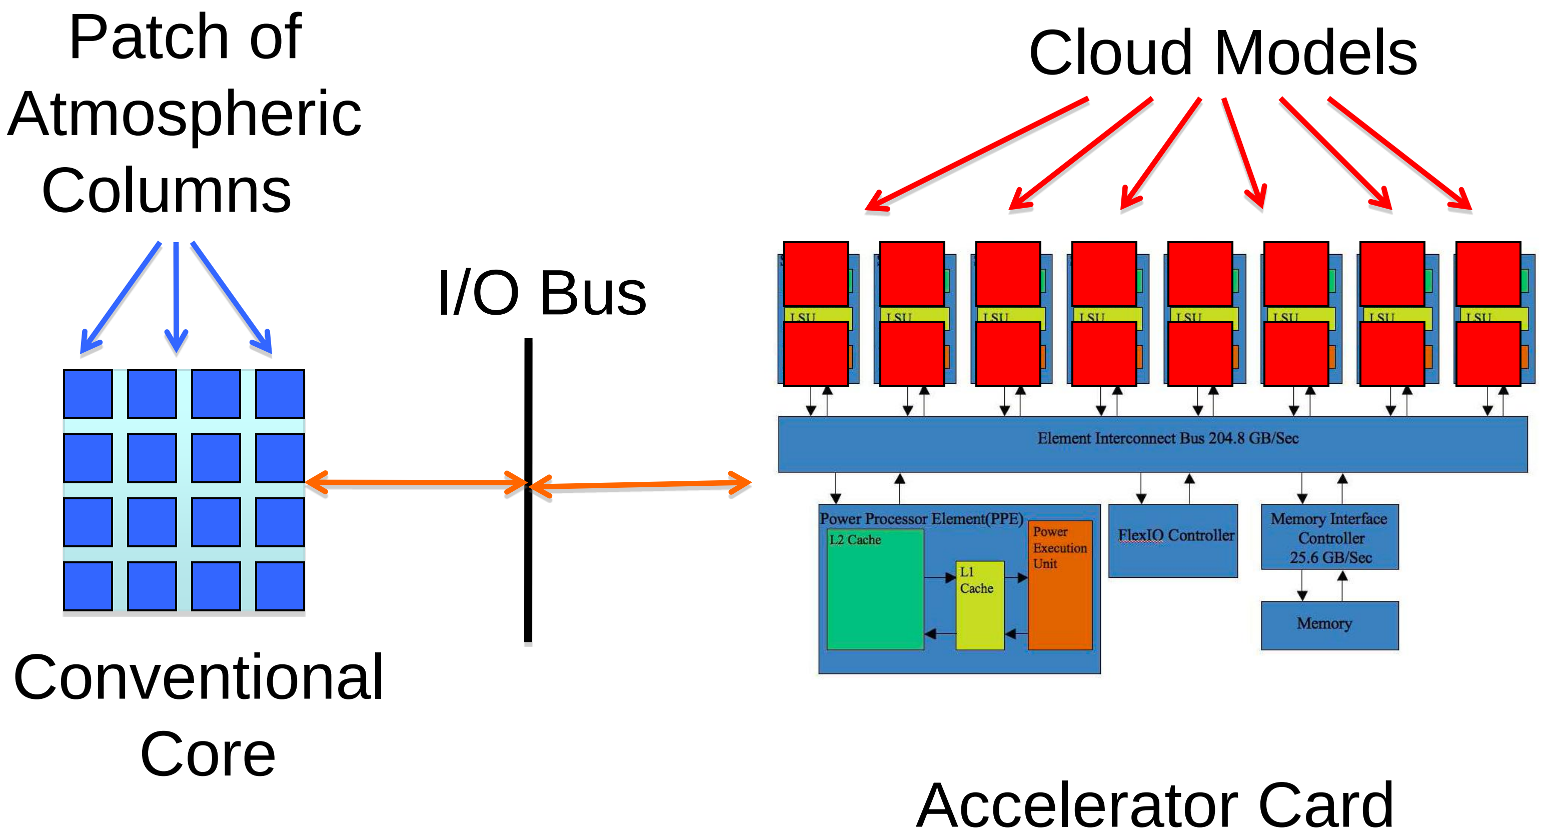

Courtesy Rich Loft, NCAR

10 Points

1. Climate change is most difficult and important problem
2. Scientific consensus: humans contribute to climate change
3. World, urgently engaged in global mitigation and regional adaptation strategies, needs detailed regional information
4. IPCC AR4: considerable uncertainty in predictions of magnitude of global change
5. … and uncertainties in regional climate are even bigger
6. Climate prediction is very computationally demanding
7. Climate spatial scales span 10 decades, but current models resolve less than 4 decades
8. Within 6 years, peak capability of 100 petaflops, 10^7 computing units
9. Parallel computation means new software, algorithms and models
10. … and new methods in workflow management, data management, and visualization

Investments Required in:

- People and infrastructure for…
  - Algorithm Research – AMR, solvers..
  - Software Engineering – DAV tools, frameworks
  - Computational Science Research – scalability, fault tolerance
- Architectural innovations including…
  - Global reduction networks
  - Robust I/O subsystems
  - Tighter integration of accelerators and µprocs
- New modeling schemes, such as…
  - pairing sub-grid-scale models with accelerators

Courtesy Rich Loft, NCAR

Working on Climate Scalability Requires Big Interdisciplinary Teams

- **Contributors:**

D. Bader (ORNL)
D. Bailey (NCAR)
C. Bitz (U Washington)
F. Bryan (NCAR)
T. Craig (NCAR)
A. St. Cyr (NCAR)
J. Dennis (NCAR)
J. Edwards (IBM)
B. Fox-Kemper (MIT,CU)
E. Hunke (LANL)
B. Kadlec (CU)
D. Ivanova (LLNL)
E. Jedlicka (ANL)
E. Jessup (CU)
R. Jacob (ANL)
P. Jones (LANL)
J, Kinter (COLA)
S. Mishra (NCAR)
S. Peacock (NCAR)
K. Lindsay (NCAR)
W. Lipscomb (LANL)
R. Loy (ANL)
J. Michalakes (NCAR)
A. Mirin (LLNL)
M. Maltrud (LANL)
J. McClean (LLNL)
R. Nair (NCAR)
M. Norman (NCSU)
T. Qian (NCAR)
C. Stan (COLA)
M. Taylor (SNL)
H. Tufo (NCAR)
M. Vertenstein (NCAR)
P. Worley (ORNL)
M. Zhang (SUNYSB)

- **Funding:**
  - DOE-BER CCPP Program Grant
    - DE-FC03-97ER62402
    - DE-PS02-07ER07-06
    - DE-FC02-07ER64340
    - B&R KP1206000
  - DOE-ASCR
    - B&R KJ0101030
  - NSF Cooperative Grant NSF01
  - NSF PetaApps Award
- **Computer Time:**
  - Blue Gene/L time:
    - NSF MRI Grant
      - NCAR
      - University of Colorado
      - IBM (SUR) program
    - BGW Consortium Days
      - IBM research (Watson)
    - LLNL
    - Stony Brook & BNL
  - CRAY XT time:
    - NICS/ORNL
    - NERSC
    - Sandia

Scale of the Enterprise

- Computational Power – **Petaflops/sec**
- Data Volumes – **Petabytes Exabytes**
- Model re-coding for massively parallel architectures → **100 dedicated person-years**
- Climate data assimilation, initialisation and reanalysis → **100 dedicated person-years**
- **Sufficient expertise** in climate and Earth system processes

Challenge will be securing critical capability in all these areas.

Is it beyond the capacity of a single country?

Is it time to tackle this together and avoid duplication of effort?

Courtesy of Julia Slingo
World Modeling Summit, May 2008, Reading, UK

Summary

- Considerable progress in weather & climate modeling over the past 45 years along with a 10^9-fold increase in computing
- Breakthroughs in the next decade will require huge increases in model resolution & complexity $\rightarrow 10^{3+}$ X increase in computing capability, along with work on entire spectrum of issues in high-end computing and model & code development
- We are not currently organized as a community to step up to this challenge
- This problem may be larger than any single nation can address … international cooperation is required to accelerate progress and productively use petascale computing for climate prediction.